

AIX 5L Version 5.3



# Technical Reference: Base Operating System and Extensions, Volume 1



AIX 5L Version 5.3



# Technical Reference: Base Operating System and Extensions, Volume 1

**Note**

Before using this information and the product it supports, read the information in Appendix C, "Notices," on page 1349.

**Fifth Edition (November 2007)**

This edition applies to AIX 5L Version 5.3 and to all subsequent releases of this product until otherwise indicated in new editions.

A reader's comment form is provided at the back of this publication. If the form has been removed, address comments to Information Development, Department 04XA-905-6C006, 11501 Burnet Road, Austin, Texas 78758-3493. To send comments electronically, use this commercial Internet address: [aix6kpub@austin.ibm.com](mailto:aix6kpub@austin.ibm.com). Any information that you supply may be used without incurring any obligation to you.

© Copyright International Business Machines Corporation 1994, 2007. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>About This Book</b> . . . . .	xxi
Highlighting . . . . .	xxi
Case-Sensitivity in AIX . . . . .	xxi
ISO 9000 . . . . .	xxi
32-Bit and 64-Bit Support for the Single UNIX Specification. . . . .	xxii
Related Publications . . . . .	xxii
<b>Base Operating System (BOS) Runtime Services (A-P)</b> . . . . .	1
a64l or l64a Subroutine . . . . .	1
abort Subroutine . . . . .	2
abs, div, labs, ldiv, imul_dbl, umul_dbl, llabs, or lldiv Subroutine . . . . .	3
access, accessx, or faccessx Subroutine . . . . .	4
acct Subroutine. . . . .	7
acl_chg or acl_fchg Subroutine . . . . .	8
acl_get or acl_fget Subroutine . . . . .	10
acl_put or acl_fput Subroutine . . . . .	12
acl_set or acl_fset Subroutine . . . . .	14
aclx_convert Subroutine . . . . .	16
aclx_get or aclx_fget Subroutine . . . . .	17
aclx_gettypeinfo Subroutine. . . . .	20
aclx_gettypes Subroutine . . . . .	21
aclx_print or aclx_printStr Subroutine . . . . .	23
aclx_put or aclx_fput Subroutine . . . . .	25
aclx_scan or aclx_scanStr Subroutine . . . . .	27
acos, acosf, or acosl Subroutine . . . . .	29
acosh, acoshf, or acoshl Subroutine . . . . .	30
addproj Subroutine . . . . .	31
addprojdb Subroutine . . . . .	32
addsys Subroutine. . . . .	33
adjtime Subroutine . . . . .	35
agg_proc_stat, agg_lpar_stat, agg_arm_stat, or free_agg_list Subroutine . . . . .	36
aio_cancel or aio_cancel64 Subroutine . . . . .	38
aio_error or aio_error64 Subroutine . . . . .	42
aio_fsync Subroutine . . . . .	45
aio_nwait Subroutine . . . . .	47
aio_nwait_timeout Subroutine . . . . .	49
aio_read or aio_read64 Subroutine . . . . .	50
aio_return or aio_return64 Subroutine . . . . .	55
aio_suspend or aio_suspend64 Subroutine . . . . .	58
aio_write or aio_write64 Subroutine . . . . .	61
alloc, dealloc, print, read_data, read_regs, symbol_addrs, write_data, and write_regs Subroutine . . . . .	66
alloca Subroutine. . . . .	68
arm_end Subroutine . . . . .	69
arm_end Dual Call Subroutine. . . . .	71
arm_getid Subroutine . . . . .	73
arm_getid Dual Call Subroutine . . . . .	75
arm_init Subroutine. . . . .	77
arm_init Dual Call Subroutine . . . . .	79
arm_start Subroutine . . . . .	81
arm_start Dual Call Subroutine . . . . .	82
arm_stop Subroutine . . . . .	84
arm_stop Dual Call Subroutine . . . . .	86
arm_update Subroutine . . . . .	88

arm_update Dual Call Subroutine . . . . .	89
asinh, asinhf, or asinhl Subroutine . . . . .	90
asinf, asinl, or asin Subroutine. . . . .	91
assert Macro . . . . .	92
atan2f, atan2l, or atan2 Subroutine . . . . .	93
atan, atanf, or atanl Subroutine . . . . .	94
atanh, atanhf, or atanhl Subroutine . . . . .	95
atof atoff Subroutine . . . . .	96
atol or atoll Subroutine . . . . .	98
audit Subroutine . . . . .	98
auditbin Subroutine . . . . .	100
auditevents Subroutine . . . . .	102
auditlog Subroutine . . . . .	104
auditobj Subroutine . . . . .	105
auditpack Subroutine. . . . .	108
auditproc Subroutine . . . . .	109
auditread, auditread_r Subroutines. . . . .	111
auditwrite Subroutine. . . . .	112
authenticate Subroutine. . . . .	113
authenticatex Subroutine . . . . .	115
basename Subroutine . . . . .	117
bcopy, bcmp, bzero or ffs Subroutine . . . . .	118
bessel: j0, j1, jn, y0, y1, or yn Subroutine . . . . .	119
bindprocessor Subroutine . . . . .	120
brk or sbrk Subroutine . . . . .	122
bsearch Subroutine . . . . .	123
btowc Subroutine . . . . .	124
buildproclist Subroutine . . . . .	125
buildtranlist or freetranlist Subroutine . . . . .	126
_check_lock Subroutine. . . . .	127
_clear_lock Subroutine . . . . .	128
cabs, cabsf, or cabsl Subroutine . . . . .	129
cacos, cacosf, or cacosl Subroutine . . . . .	129
cacosh, cacoshf, or cacoshl Subroutines . . . . .	130
carg, cargf, or cargl Subroutine . . . . .	131
casin, casinf, or casinl Subroutine . . . . .	131
casinh, casinhf, or casinhl Subroutine . . . . .	132
catan, catanf, or catanl Subroutine. . . . .	132
catanh, catanhf, or catanhl Subroutine . . . . .	133
catclose Subroutine . . . . .	134
catgets Subroutine . . . . .	135
catopen Subroutine . . . . .	136
cbrtf, cbrtl, or cbrt Subroutine . . . . .	137
ccos, ccosf, or ccosl Subroutine. . . . .	138
ccosh, ccoshf, or ccoshl Subroutine . . . . .	139
ccsidtoocs or cstoccsid Subroutine . . . . .	139
ceil, ceilf, ceill, ceild32, ceild64, and ceild128 Subroutines . . . . .	140
cexp, cexpf, or cexpl Subroutine . . . . .	141
cfgetspeed, cfsetspeed, cfgetispeed, or cfsetispeed Subroutine . . . . .	142
chacl or fchacl Subroutine . . . . .	144
chdir Subroutine . . . . .	147
chmod or fchmod Subroutine. . . . .	148
chown, fchown, lchown, chownx, or fchownx Subroutine. . . . .	151
chpass Subroutine . . . . .	154
chpassx Subroutine . . . . .	156
chprojattr Subroutine. . . . .	158

chprojattdb Subroutine . . . . .	159
chroot Subroutine . . . . .	160
chsys Subroutine . . . . .	162
cimag, cimagn, or cimagl Subroutine . . . . .	163
ckuseracct Subroutine . . . . .	164
ckuserID Subroutine . . . . .	166
class, _class, finite, isnan, or unordered Subroutines . . . . .	167
clock Subroutine . . . . .	169
clock_getcpuclockid Subroutine . . . . .	169
clock_getres, clock_gettime, and clock_settime Subroutine . . . . .	170
clock_nanosleep Subroutine . . . . .	172
clog, clogf, or clogl Subroutine . . . . .	174
close Subroutine . . . . .	175
compare_and_swap Subroutine . . . . .	176
compile, step, or advance Subroutine . . . . .	177
confstr Subroutine . . . . .	181
conj, conjf, or conjl Subroutine . . . . .	182
conv Subroutines . . . . .	183
copysign, copysignf, copysignl, copysignd32, copysignd64, and copysignd128 Subroutines . . . . .	185
coredump Subroutine . . . . .	186
cosf, cosl, or cos Subroutine . . . . .	187
cosh, coshf, or coshl Subroutine . . . . .	188
cpow, cpowf, or cpowl Subroutine . . . . .	189
cproj, cprojf, or cprojl Subroutine . . . . .	190
creal, crealf, or creall Subroutine . . . . .	190
crypt, encrypt, or setkey Subroutine . . . . .	191
csid Subroutine . . . . .	193
csin, csinf, or csinl Subroutine . . . . .	193
csinh, csinhf, or csinhl Subroutine . . . . .	194
csqrt, csqrtf, or csqrtl Subroutine . . . . .	195
CT_HOOKx and CT_GEN macros . . . . .	195
CT_TRCON macro . . . . .	197
ctan, ctanf, or ctanl Subroutine . . . . .	198
ctanh, ctanhf, or ctanhl Subroutine . . . . .	198
ctermid Subroutine . . . . .	199
ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine . . . . .	200
ctime64, localtime64, gmtime64, mktime64, difftime64, or asctime64 Subroutine . . . . .	202
ctime64_r, localtime64_r, gmtime64_r, or asctime64_r Subroutine . . . . .	204
ctime_r, localtime_r, gmtime_r, or asctime_r Subroutine . . . . .	206
ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines . . . . .	208
cuserid Subroutine . . . . .	210
defsys Subroutine . . . . .	211
delssys Subroutine . . . . .	212
dirname Subroutine . . . . .	213
disclaim and disclaim64 Subroutines . . . . .	214
dlclose Subroutine . . . . .	215
dLError Subroutine . . . . .	216
dlopen Subroutine . . . . .	217
dlsym Subroutine . . . . .	219
drand48, erand48, jrand48, lcong48, lrand48, mrand48, nrand48, seed48, or srand48 Subroutine . . . . .	221
drem Subroutine . . . . .	223
_end, _etext, or _edata Identifier . . . . .	224
ecvt, fcvt, or gcvrt Subroutine . . . . .	225
EnableCriticalSections, BeginCriticalSection, and EndCriticalSection Subroutine . . . . .	226
erf, erff, or erfl Subroutine . . . . .	227

erfc, erfcf, or erfcl Subroutine . . . . .	228
errlog Subroutine . . . . .	229
errlog_close Subroutine . . . . .	231
errlog_find_first, errlog_find_next, and errlog_find_sequence Subroutines . . . . .	232
errlog_open Subroutine . . . . .	234
errlog_set_direction Subroutine . . . . .	235
errlog_write Subroutine . . . . .	235
exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine . . . . .	236
exit, atexit, unatexit, _exit, or _Exit Subroutine . . . . .	243
exp, expf, or expl Subroutine . . . . .	245
exp2, exp2f, or exp2l Subroutine . . . . .	247
expm1, expm1f, or expm1l Subroutine . . . . .	248
fabsf, fabsl, fabs, fabsd32, fabsd64, and fabsd128 Subroutines . . . . .	249
fattach Subroutine . . . . .	250
fchdir Subroutine . . . . .	251
fclear or fclear64 Subroutine . . . . .	252
fclose or fflush Subroutine . . . . .	253
fcntl, dup, or dup2 Subroutine . . . . .	255
fdetach Subroutine . . . . .	261
fdim, fdimf, fdiml, fdimd32, fdimd64, and fdimd128 Subroutines . . . . .	262
fe_dec_getround and fe_dec_setround Subroutines . . . . .	263
feclearexcept Subroutine . . . . .	265
fegetenv or fesetenv Subroutine . . . . .	265
fegetexceptflag or fesetexceptflag Subroutine . . . . .	266
fegetround or fesetround Subroutine . . . . .	267
feholdexcept Subroutine . . . . .	267
fence Subroutine . . . . .	268
feof, ferror, clearerr, or fileno Macro . . . . .	270
feraiseexcept Subroutine . . . . .	270
fetch_and_add Subroutine . . . . .	271
fetch_and_and or fetch_and_or Subroutine . . . . .	272
fetestexcept Subroutine . . . . .	273
feupdateenv Subroutine . . . . .	274
finfo or ffinfo Subroutine . . . . .	274
flockfile, ftrylockfile, funlockfile Subroutine . . . . .	275
floor, floorf, floorl, floord32, floord64, floord128, nearest, trunc, itrunc, and uitrunc Subroutines . . . . .	277
fma, fmaf, or fmal Subroutine . . . . .	279
fmax, fmaxf, fmaxl, fmaxd32, fmaxd64, and fmaxd128 Subroutines . . . . .	280
fminf, fminl, fmind32, fmind64, and fmind128 Subroutines . . . . .	281
fmod, fmodf, or fmodl Subroutine . . . . .	282
fmtmsg Subroutine . . . . .	283
fnmatch Subroutine . . . . .	286
fopen, fopen64, freopen, freopen64 or fdopen Subroutine . . . . .	288
fork, f_fork, or vfork Subroutine . . . . .	291
fp_any_enable, fp_is_enabled, fp_enable_all, fp_enable, fp_disable_all, or fp_disable Subroutine . . . . .	294
fp_clr_flag, fp_set_flag, fp_read_flag, or fp_swap_flag Subroutine . . . . .	295
fp_cpusync Subroutine . . . . .	297
fp_flush_impresise Subroutine . . . . .	298
fp_invalid_op, fp_divbyzero, fp_overflow, fp_underflow, fp_inexact, fp_any_xcp Subroutine . . . . .	299
fp_iop_snan, fp_iop_infsinf, fp_iop_infdfinf, fp_iop_zrdzr, fp_iop_infmzr, fp_iop_invcmp, fp_iop_sqrt, fp_iop_convert, or fp_iop_vxsoft Subroutines . . . . .	300
fp_raise_xcp Subroutine . . . . .	301
fp_read_rnd or fp_swap_rnd Subroutine . . . . .	302
fp_sh_info, fp_sh_trap_info, or fp_sh_set_stat Subroutine . . . . .	303
fp_trap Subroutine . . . . .	306
fp_trapstate Subroutine . . . . .	308



fpclassify Macro . . . . .	309
fread or fwrite Subroutine . . . . .	310
freehostent Subroutine . . . . .	312
freelmb Subroutine . . . . .	313
frevolve Subroutine . . . . .	313
frexp32, frexp64, and frexp128 Subroutines . . . . .	314
frexpf, frexpl, or frexp Subroutine . . . . .	315
fscntl Subroutine . . . . .	316
fseek, fseeko, fseeko64, rewind, ftell, ftello, ftello64, fgetpos, fgetpos64, fsetpos, or fsetpos64 Subroutine. . . . .	318
fsync or fsync_range Subroutine . . . . .	321
ftok Subroutine . . . . .	322
ftw or ftw64 Subroutine . . . . .	324
fwide Subroutine . . . . .	326
fwprintf, wprintf, swprintf Subroutines . . . . .	327
fwscanf, wscanf, swscanf Subroutines . . . . .	331
gai_strerror Subroutine . . . . .	336
gamma Subroutine . . . . .	336
gencore or coredump Subroutine . . . . .	337
genpagvalue Subroutine . . . . .	339
get_malloc_log Subroutine. . . . .	341
get_malloc_log_live Subroutine . . . . .	341
get_speed, set_speed, or reset_speed Subroutines . . . . .	342
getargs Subroutine . . . . .	343
getauditostattr, IDtohost, hosttoID, nexthost or putauditostattr Subroutine . . . . .	345
getauthdb or getauthdb_r Subroutine . . . . .	346
getc, getchar, fgetc, or getw Subroutine . . . . .	347
getc_unlocked, getchar_unlocked, putc_unlocked, putchar_unlocked Subroutines . . . . .	349
getconfattr or putconfattr Subroutine . . . . .	350
getconfattr_s Subroutine . . . . .	355
getcontext or setcontext Subroutine . . . . .	357
getcwd Subroutine . . . . .	358
getdate Subroutine . . . . .	360
getdtablesize Subroutine . . . . .	363
getea Subroutine . . . . .	363
getenv Subroutine. . . . .	365
getenvs Subroutine . . . . .	365
getfilehdr Subroutine . . . . .	367
getfirstprojdb Subroutine . . . . .	368
getfsent, getfsspec, getfsfile, getfstype, setfsent, or endfsent Subroutine . . . . .	369
getgid, getegid or gegidx Subroutine . . . . .	370
getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine . . . . .	371
getgrgid_r Subroutine . . . . .	373
getgrnam_r Subroutine . . . . .	374
getgroupattr, IDtogroup, nextgroup, or putgroupattr Subroutine . . . . .	375
getgroupattr_s Subroutine . . . . .	378
getgroups Subroutine . . . . .	383
getgrpaclattr, nextgrpacl, or putgrpaclattr Subroutine . . . . .	384
getgrset Subroutine . . . . .	385
getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine . . . . .	386
getiopri Subroutine . . . . .	389
getipnodebyaddr Subroutine . . . . .	390
getipnodebyname Subroutine . . . . .	391
getlogin Subroutine . . . . .	393
getlogin_r Subroutine . . . . .	395
getnextprojdb Subroutine . . . . .	396

getopt Subroutine . . . . .	397
getpagesize Subroutine . . . . .	399
getpaginfo Subroutine . . . . .	400
getpagvalue or getpagvalue64 Subroutine . . . . .	401
getpass Subroutine . . . . .	402
getpcred Subroutine . . . . .	403
getpeereid Subroutine . . . . .	405
getpenv Subroutine . . . . .	405
getpgid Subroutine . . . . .	407
getpid, getpgrp, or getppid Subroutine . . . . .	407
getportattr or putportattr Subroutine . . . . .	408
getpri Subroutine . . . . .	411
getpriority, setpriority, or nice Subroutine . . . . .	412
getproclist, getparlist, or getarmlist Subroutine . . . . .	414
getprocs Subroutine . . . . .	415
getproj Subroutine . . . . .	418
getprojdb Subroutine . . . . .	419
getprojs Subroutine . . . . .	420
getpw Subroutine . . . . .	421
getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine . . . . .	422
getrlimit, getrlimit64, setrlimit, setrlimit64, or vlimit Subroutine . . . . .	424
getrpcnt, getrpcbyname, getrpcbynumber, setrpcnt, or endrpcnt Subroutine . . . . .	427
getrusage, getrusage64, times, or vtimes Subroutine . . . . .	428
getroleattr, nextrole or putroleattr Subroutine . . . . .	431
gets or fgets Subroutine . . . . .	434
getfsent_r, getfsspec_r, getfsfile_r, getfstype_r, setfsent_r, or endfsent_r Subroutine . . . . .	435
getsid Subroutine . . . . .	436
getssys Subroutine . . . . .	437
getsubopt Subroutine . . . . .	438
getsubsvr Subroutine . . . . .	439
gettcattr or puttcattr Subroutine . . . . .	440
getthrd Subroutine . . . . .	443
gettimeofday, settimeofday, or ftime Subroutine . . . . .	445
gettimer, settimer, restimer, stime, or time Subroutine . . . . .	446
gettimerid Subroutine . . . . .	449
getttyent, gettyent, setttyent, or endttyent Subroutine . . . . .	450
getuid, geteuid, or getuidx Subroutine . . . . .	452
getuinfo Subroutine . . . . .	453
getuinfox Subroutine . . . . .	453
getuserattr, IDtouser, nextuser, or putuserattr Subroutine . . . . .	454
getuserattrx Subroutine . . . . .	460
GetUserAuths Subroutine . . . . .	467
getuserpw, putuserpw, or putuserpwhist Subroutine . . . . .	468
getuserpwx Subroutine . . . . .	470
getusraclattr, nextusracl or putusraclattr Subroutine . . . . .	472
getutent, getutid, getutline, pututline, setutent, endutent, or utmpname Subroutine . . . . .	474
getvfsent, getvfsbytype, getvfsbyname, getvfsbyflag, setvfsent, or endvfsent Subroutine . . . . .	476
getwc, fgetwc, or getwchar Subroutine . . . . .	477
getwd Subroutine . . . . .	479
getws or fgetws Subroutine . . . . .	480
glob Subroutine . . . . .	482
globfree Subroutine . . . . .	485
grantpt Subroutine . . . . .	485
HBA_CloseAdapter Subroutine . . . . .	486
HBA_FreeLibrary Subroutine . . . . .	487

HBA_GetAdapterAttributes, HBA_GetPortAttributes, HBA_GetDiscoveredPortAttributes, HBA_GetPortAttributesByWWN Subroutine . . . . .	487
HBA_GetAdapterName Subroutine . . . . .	490
HBA_GetEventBuffer Subroutine . . . . .	491
HBA_GetFC4Statistics Subroutine . . . . .	492
HBA_GetFcpPersistentBinding Subroutine . . . . .	493
HBA_GetFCPStatistics Subroutine . . . . .	494
HBA_GetFcpTargetMappingV2 Subroutine . . . . .	495
HBA_GetFcpTargetMapping Subroutine . . . . .	496
HBA_GetNumberOfAdapters Subroutine . . . . .	497
HBA_GetPersistentBindingV2 Subroutine . . . . .	498
HBA_GetPortStatistics Subroutine . . . . .	499
HBA_GetRNIDMgmtInfo Subroutine . . . . .	500
HBA_GetVersion Subroutine . . . . .	501
HBA_LoadLibrary Subroutine . . . . .	502
HBA_OpenAdapter Subroutine . . . . .	503
HBA_OpenAdapterByWWN Subroutine . . . . .	503
HBA_RefreshInformation Subroutine . . . . .	504
HBA_ScsiInquiryV2 Subroutine . . . . .	505
HBA_ScsiReadCapacityV2 Subroutine . . . . .	507
HBA_ScsiReportLunsV2 Subroutine . . . . .	509
HBA_SendCTPassThru Subroutine . . . . .	510
HBA_SendCTPassThruV2 Subroutine . . . . .	511
HBA_SendReadCapacity Subroutine . . . . .	512
HBA_SendReportLUNs Subroutine . . . . .	513
HBA_SendRLS Subroutine . . . . .	515
HBA_SendRNID Subroutine . . . . .	516
HBA_SendRNIDV2 Subroutine . . . . .	517
HBA_SendRPL Subroutine . . . . .	519
HBA_SendRPS Subroutine . . . . .	520
HBA_SendScsiInquiry Subroutine . . . . .	521
HBA_SetRNIDMgmtInfo Subroutine . . . . .	523
hpmInit, f_hpminit, hpmStart, f_hpmstart, hpmStop, f_hpmstop, hpmTstart, f_hpmtstart, hpmTstop, f_hpmtstop, hpmGetTimeAndCounters, f_hpmgettimeandcounters, hpmGetCounters, f_hpmgetcounters, hpmTerminate, and f_hpmtterminate Subroutine . . . . .	524
hsearch, hcreate, or hdestroy Subroutine . . . . .	527
hypot, hypotf, or hypotl Subroutine . . . . .	528
iconv Subroutine . . . . .	529
iconv_close Subroutine . . . . .	531
iconv_open Subroutine . . . . .	532
ilogbf, ilogbl, or ilogb Subroutine . . . . .	534
imaxabs Subroutine . . . . .	535
imaxdiv Subroutine . . . . .	535
IMAIXMapping Subroutine . . . . .	536
IMAuxCreate Callback Subroutine . . . . .	537
IMAuxDestroy Callback Subroutine . . . . .	537
IMAuxDraw Callback Subroutine . . . . .	538
IMAuxHide Callback Subroutine . . . . .	539
IMBeep Callback Subroutine . . . . .	539
IMClose Subroutine . . . . .	540
IMCreate Subroutine . . . . .	541
IMDestroy Subroutine . . . . .	541
IMFilter Subroutine . . . . .	542
IMFreeKeymap Subroutine . . . . .	543
IMIndicatorDraw Callback Subroutine . . . . .	543
IMIndicatorHide Callback Subroutine . . . . .	544

IMInitialize Subroutine . . . . .	545
IMInitializeKeymap Subroutine . . . . .	546
IMIoctl Subroutine . . . . .	547
IMLookupString Subroutine . . . . .	549
IMProcess Subroutine . . . . .	550
IMProcessAuxiliary Subroutine . . . . .	551
IMQueryLanguage Subroutine . . . . .	552
IMSimpleMapping Subroutine . . . . .	553
IMTextCursor Callback Subroutine . . . . .	554
IMTextDraw Callback Subroutine . . . . .	555
IMTextHide Callback Subroutine . . . . .	555
IMTextStart Callback Subroutine . . . . .	556
inet_aton Subroutine . . . . .	557
initgroups Subroutine. . . . .	558
initialize Subroutine . . . . .	558
insque or remque Subroutine. . . . .	559
install_lwcf_handler Subroutine . . . . .	560
ioctl, ioctlx, ioctl32, or ioctl32x Subroutine . . . . .	561
isblank Subroutine. . . . .	564
isendwin Subroutine . . . . .	565
isfinite Macro . . . . .	566
isgreater Macro. . . . .	566
isgreaterequal Subroutine . . . . .	567
isinf Subroutine. . . . .	568
isless Macro . . . . .	568
islessequal Macro . . . . .	569
islessgreater Macro . . . . .	570
isnormal Macro . . . . .	570
isunordered Macro . . . . .	571
iswalnum, iswalph, iswcntrl, iswdigit, iswgraph, iswlower, iswprint, iswpunct, iswspace, iswupper, or iswxdigit Subroutine . . . . .	572
iswblank Subroutine . . . . .	573
iswctype or is_wctype Subroutine . . . . .	574
jcode Subroutines . . . . .	575
Japanese conv Subroutines . . . . .	576
Japanese ctype Subroutines . . . . .	578
kill or killpg Subroutine . . . . .	580
kleenup Subroutine . . . . .	581
knlist Subroutine . . . . .	582
kpidxstate Subroutine . . . . .	584
_lazySetErrorHandler Subroutine . . . . .	584
l3tol or ltol3 Subroutine . . . . .	586
l64a_r Subroutine . . . . .	587
LAPI_Addr_get Subroutine . . . . .	588
LAPI_Addr_set Subroutine . . . . .	589
LAPI_Address Subroutine. . . . .	591
LAPI_Address_init Subroutine . . . . .	592
LAPI_Address_init64 Subroutine . . . . .	594
LAPI_Amsend Subroutine. . . . .	596
LAPI_Amsendv Subroutine . . . . .	601
LAPI_Fence Subroutine . . . . .	609
LAPI_Get Subroutine . . . . .	610
LAPI_Getcptr Subroutine . . . . .	613
LAPI_Getv Subroutine . . . . .	614
LAPI_Gfence Subroutine . . . . .	618
LAPI_Init Subroutine. . . . .	619

LAPI_Msg_string Subroutine . . . . .	624
LAPI_Msgpoll Subroutine . . . . .	626
LAPI_Nopoll_wait Subroutine . . . . .	628
LAPI_Probe Subroutine . . . . .	629
LAPI_Purge_totask Subroutine . . . . .	630
LAPI_Put Subroutine . . . . .	632
LAPI_Putv Subroutine . . . . .	634
LAPI_Qenv Subroutine . . . . .	638
LAPI_Resume_totask Subroutine . . . . .	641
LAPI_Rmw Subroutine . . . . .	642
LAPI_Rmw64 Subroutine . . . . .	646
LAPI_Senv Subroutine . . . . .	650
LAPI_Setcncr Subroutine . . . . .	652
LAPI_Setcncr_wstatus Subroutine . . . . .	654
LAPI_Term Subroutine . . . . .	655
LAPI_Util Subroutine . . . . .	657
LAPI_Waitcncr Subroutine . . . . .	668
LAPI_Xfer Subroutine . . . . .	670
layout_object_create Subroutine . . . . .	683
layout_object_editshape or wcslayout_object_editshape Subroutine . . . . .	684
layout_object_getvalue Subroutine . . . . .	687
layout_object_setvalue Subroutine . . . . .	689
layout_object_shapeboxchars Subroutine . . . . .	691
layout_object_transform or wcslayout_object_transform Subroutine . . . . .	692
layout_object_free Subroutine . . . . .	695
ldahread Subroutine . . . . .	696
ldclose or ldaclose Subroutine . . . . .	697
ldexpd32, ldexpd64, and ldexpd128 Subroutines . . . . .	698
ldexp, ldexpf, or ldexpl Subroutine . . . . .	699
ldfread Subroutine . . . . .	701
ldgetname Subroutine . . . . .	702
ldlread, ldlnit, or ldlnitem Subroutine . . . . .	704
ldlseek or ldlnseek Subroutine . . . . .	705
ldohseek Subroutine . . . . .	707
ldopen or ldaopen Subroutine . . . . .	707
ldrseek or ldnrseek Subroutine . . . . .	709
ldshread or ldnsbread Subroutine . . . . .	710
ldsseek or ldnsseek Subroutine . . . . .	712
ldtbindx Subroutine . . . . .	713
ldtbread Subroutine . . . . .	714
ldtbseek Subroutine . . . . .	715
lgamma, lgammaf, or lgammal Subroutine . . . . .	716
lineout Subroutine . . . . .	717
link Subroutine . . . . .	718
lio_listio or lio_listio64 Subroutine . . . . .	720
listea Subroutine . . . . .	724
llrint, llrintf, llrintl, llrintd32, llrintd64, and llrintd128 Subroutines . . . . .	725
llround, llroundf, llroundl, llroundd32, llroundd64, and llroundd128 Subroutines . . . . .	726
load and loadAndInit Subroutines . . . . .	727
loadbind Subroutine . . . . .	731
loadquery Subroutine . . . . .	733
localeconv Subroutine . . . . .	735
lockfx, lockf, flock, or lockf64 Subroutine . . . . .	740
log10, log10f, or log10l Subroutine . . . . .	743
log1p, log1pf, or log1pl Subroutine . . . . .	744
log2, log2f, or log2l Subroutine . . . . .	745

logbd32, logbd64, and logbd128 Subroutines . . . . .	746
ilogbd32, ilogbd64, and ilogbd128 Subroutines . . . . .	747
logbf, logbl, or logb Subroutine . . . . .	748
log, logf, or logl Subroutine . . . . .	749
loginfailed Subroutine . . . . .	750
loginrestrictions Subroutine . . . . .	752
loginrestrictionsx Subroutine . . . . .	754
loginsuccess Subroutine . . . . .	757
lpar_get_info Subroutine . . . . .	759
lpar_set_resources Subroutine . . . . .	760
lrint, lrintf, lrintl, lrintd32, lrintd64, and lrintd128 Subroutines . . . . .	761
lround, lroundf, lroundl, lroundd32, lroundd64, and lroundd128 Subroutines. . . . .	763
lsearch or lfind Subroutine . . . . .	764
lseek, llseek or lseek64 Subroutine . . . . .	765
lvm_querylv Subroutine . . . . .	766
lvm_querypv Subroutine . . . . .	770
lvm_queryvg Subroutine . . . . .	774
lvm_queryvgs Subroutine . . . . .	777
malloc, free, realloc, calloc, mallopt, mallinfo, mallinfo_heap, alloca, valloc, or posix_memalign Subroutine. . . . .	778
madd, msub, mult, mdiv, pow, gcd, invert, rpow, msqrt, mcmp, move, min, omin, fmin, m_in, mout, omout, fmout, m_out, sdiv, or itom Subroutine. . . . .	785
madvise Subroutine . . . . .	787
makecontext or swapcontext Subroutine . . . . .	788
matherr Subroutine . . . . .	789
MatchAllAuths, MatchAnyAuths, MatchAllAuthsList, or MatchAnyAuthsList Subroutine . . . . .	790
mblen Subroutine . . . . .	791
mbrlen Subroutine. . . . .	792
mbrtowc Subroutine . . . . .	793
mbsadvance Subroutine . . . . .	794
mbscat, mbscmp, or mbscpy Subroutine . . . . .	795
mbschr Subroutine . . . . .	796
mbsinit Subroutine . . . . .	797
mbsinvalid Subroutine . . . . .	798
mbslen Subroutine . . . . .	798
mbsncat, mbsncmp, or mbsncpy Subroutine . . . . .	799
mbspbrk Subroutine . . . . .	800
mbsrchr Subroutine . . . . .	801
mbsrtowcs Subroutine . . . . .	802
mbstomb Subroutine . . . . .	803
mbstowcs Subroutine . . . . .	804
mbswidth Subroutine. . . . .	805
mbtowc Subroutine . . . . .	805
memccpy, memchr, memcmp, memcpy, memset or memmove Subroutine . . . . .	807
mincore Subroutine . . . . .	808
MIO_ao_read64 Subroutine . . . . .	809
MIO_ao_suspend64 Subroutine . . . . .	810
MIO_ao_write64 Subroutine . . . . .	811
MIO_close Subroutine . . . . .	812
MIO_fcntl Subroutine. . . . .	816
MIO_ffinfo Subroutine . . . . .	817
MIO_fstat64 Subroutine. . . . .	818
MIO_fsync Subroutine . . . . .	819
MIO_ftruncate64 Subroutine . . . . .	819
MIO_lio_listio64 Subroutine . . . . .	820
MIO_lseek64 Subroutine . . . . .	821

MIO_open64 Subroutine . . . . .	822
MIO_open Subroutine . . . . .	828
MIO_read Subroutine . . . . .	833
MIO_write Subroutine . . . . .	834
mkdir Subroutine . . . . .	835
mknod or mkfifo Subroutine . . . . .	837
mktemp or mkstemp Subroutine . . . . .	838
mlock and munlock Subroutine . . . . .	840
mlockall and munlockall Subroutine . . . . .	841
mmap or mmap64 Subroutine . . . . .	843
mntctl Subroutine . . . . .	847
modf, modff, modfl, modfd32, modfd64, and modfd128 Subroutines . . . . .	848
moncontrol Subroutine . . . . .	849
monitor Subroutine . . . . .	851
monstartup Subroutine . . . . .	856
mprotect Subroutine . . . . .	859
mq_close Subroutine . . . . .	861
mq_getattr Subroutine . . . . .	862
mq_notify Subroutine . . . . .	863
mq_open Subroutine . . . . .	864
mq_receive Subroutine . . . . .	866
mq_send Subroutine . . . . .	868
mq_setattr Subroutine . . . . .	869
mq_receive, mq_timedreceive Subroutine . . . . .	870
mq_send, mq_timedsend Subroutine . . . . .	872
mq_unlink Subroutine . . . . .	874
msem_init Subroutine . . . . .	875
msem_lock Subroutine . . . . .	876
msem_remove Subroutine . . . . .	877
msem_unlock Subroutine . . . . .	878
msgctl Subroutine . . . . .	879
msgget Subroutine . . . . .	881
msgrcv Subroutine . . . . .	882
msgsnd Subroutine . . . . .	885
msgxrcv Subroutine . . . . .	887
msleep Subroutine . . . . .	889
msync Subroutine . . . . .	890
mt_trce Subroutine . . . . .	891
munmap Subroutine . . . . .	893
mwakeup Subroutine . . . . .	894
nan, nanf, nanl, nand32, nand64, and nand128 Subroutines . . . . .	895
nanosleep Subroutine . . . . .	896
nearbyint, nearbyintf, nearbyintl, nearbyintd32, nearbyintd64, and nearbyintd128 Subroutines . . . . .	897
nextafter, nextafterf, nextafterl, nexttoward, nexttowardf, or nexttowardl Subroutine . . . . .	898
nextafterd32, nextafterd64, nextafterd128, nexttowardd32, nexttowardd64, and nexttowardd128 Subroutines . . . . .	900
newpass Subroutine . . . . .	901
newpassx Subroutine . . . . .	903
nftw or nftw64 Subroutine . . . . .	905
nl_langinfo Subroutine . . . . .	907
nlist, nlist64 Subroutine . . . . .	909
ns_addr Subroutine . . . . .	910
ns_ntoa Subroutine . . . . .	911
odm_add_obj Subroutine . . . . .	912
odm_change_obj Subroutine . . . . .	913
odm_close_class Subroutine . . . . .	914

odm_create_class Subroutine . . . . .	915
odm_err_msg Subroutine . . . . .	916
odm_free_list Subroutine . . . . .	917
odm_get_by_id Subroutine . . . . .	918
odm_get_list Subroutine . . . . .	919
odm_get_obj, odm_get_first, or odm_get_next Subroutine . . . . .	921
odm_initialize Subroutine . . . . .	923
odm_lock Subroutine. . . . .	923
odm_mount_class Subroutine . . . . .	925
odm_open_class or odm_open_class_rdnly Subroutine . . . . .	926
odm_rm_by_id Subroutine. . . . .	927
odm_rm_class Subroutine . . . . .	928
odm_rm_obj Subroutine . . . . .	929
odm_run_method Subroutine. . . . .	931
odm_set_path Subroutine . . . . .	932
odm_set_perms Subroutine . . . . .	933
odm_terminate Subroutine. . . . .	933
odm_unlock Subroutine. . . . .	934
open, openx, open64, creat, or creat64 Subroutine. . . . .	935
opendir, readdir, telldir, seekdir, rewinddir, closedir, opendir64, readdir64, telldir64, seekdir64, rewinddir64, or closedir64 Subroutine . . . . .	943
pam_acct_mgmt Subroutine . . . . .	946
pam_authenticate Subroutine . . . . .	947
pam_chauthtok Subroutine . . . . .	949
pam_close_session Subroutine . . . . .	950
pam_end Subroutine . . . . .	951
pam_get_data Subroutine . . . . .	952
pam_get_item Subroutine . . . . .	953
pam_get_user Subroutine . . . . .	954
pam_getenv Subroutine. . . . .	955
pam_getenvlist Subroutine. . . . .	956
pam_open_session Subroutine . . . . .	957
pam_putenv Subroutine. . . . .	958
pam_set_data Subroutine . . . . .	959
pam_set_item Subroutine . . . . .	960
pam_setcred Subroutine . . . . .	961
pam_sm_acct_mgmt Subroutine . . . . .	963
pam_sm_authenticate Subroutine . . . . .	964
pam_sm_chauthtok Subroutine . . . . .	965
pam_sm_close_session Subroutine . . . . .	967
pam_sm_open_session Subroutine . . . . .	968
pam_sm_setcred Subroutine . . . . .	969
pam_start Subroutine . . . . .	970
pam_strerror Subroutine . . . . .	973
passwdexpired Subroutine. . . . .	973
passwdexpiredx Subroutine . . . . .	974
passwdpolicy Subroutine . . . . .	976
passwdstrength Subroutine . . . . .	978
pathconf or fpathconf Subroutine . . . . .	979
pause Subroutine . . . . .	982
pcap_close Subroutine . . . . .	982
pcap_compile Subroutine . . . . .	983
pcap_datalink Subroutine . . . . .	984
pcap_dispatch Subroutine . . . . .	984
pcap_dump Subroutine . . . . .	986
pcap_dump_close Subroutine . . . . .	986



pcap_dump_open Subroutine . . . . .	987
pcap_file Subroutine . . . . .	988
pcap_fileno Subroutine . . . . .	988
pcap_geterr Subroutine . . . . .	989
pcap_is_swapped Subroutine . . . . .	990
pcap_lookupdev Subroutine . . . . .	990
pcap_lookupnet Subroutine . . . . .	991
pcap_loop Subroutine . . . . .	992
pcap_major_version Subroutine . . . . .	993
pcap_minor_version Subroutine . . . . .	994
pcap_next Subroutine . . . . .	995
pcap_open_live Subroutine . . . . .	995
pcap_open_live_sb Subroutine . . . . .	996
pcap_open_offline Subroutine . . . . .	997
pcap_perror Subroutine . . . . .	998
pcap_setfilter Subroutine . . . . .	999
pcap_snapshot Subroutine . . . . .	1000
pcap_stats Subroutine . . . . .	1000
pcap_strerror Subroutine . . . . .	1001
pclose Subroutine . . . . .	1002
perfstat_cpu Subroutine . . . . .	1003
perfstat_cpu_total Subroutine . . . . .	1004
perfstat_disk Subroutine . . . . .	1006
perfstat_diskadapter Subroutine . . . . .	1008
perfstat_diskpath Subroutine . . . . .	1009
perfstat_disk_total Subroutine . . . . .	1011
perfstat_memory_total Subroutine . . . . .	1013
perfstat_netbuffer Subroutine . . . . .	1014
perfstat_netinterface Subroutine . . . . .	1015
perfstat_netinterface_total Subroutine . . . . .	1017
perfstat_pagingspace Subroutine . . . . .	1018
perfstat_partial_reset Subroutine . . . . .	1019
perfstat_partition_total Subroutine . . . . .	1021
perfstat_protocol Subroutine . . . . .	1022
perfstat_reset Subroutine . . . . .	1023
perror Subroutine . . . . .	1024
pipe Subroutine . . . . .	1025
plock Subroutine . . . . .	1026
pm_cycles Subroutine . . . . .	1027
pm_delete_program Subroutine . . . . .	1028
pm_delete_program_group Subroutine . . . . .	1029
pm_delete_program_mygroup Subroutine . . . . .	1030
pm_delete_program_mythread Subroutine . . . . .	1031
pm_delete_program_pgroup Subroutine . . . . .	1032
pm_delete_program_pthread Subroutine . . . . .	1033
pm_delete_program_thread Subroutine . . . . .	1034
pm_error Subroutine . . . . .	1035
pm_get_data, pm_get_tdata, pm_get_Tdata, pm_get_data_cpu, pm_get_tdata_cpu, pm_get_Tdata_cpu, pm_get_data_lcpu, pm_get_tdata_lcpu and pm_get_Tdata_lcpu Subroutine . . . . .	1036
pm_get_data_group, pm_get_tdata_group and pm_get_Tdata_group Subroutine . . . . .	1039
pm_get_data_group_mx and pm_get_tdata_group_mx Subroutine . . . . .	1041
pm_get_data_mx, pm_get_tdata_mx, pm_get_data_cpu_mx, pm_get_tdata_cpu_mx, pm_get_data_lcpu_mx and pm_get_tdata_lcpu_mx Subroutine . . . . .	1042
pm_get_data_mygroup, pm_get_tdata_mygroup or pm_get_Tdata_mygroup Subroutine . . . . .	1044
pm_get_data_mygroup_mx or pm_get_tdata_mygroup_mx Subroutine . . . . .	1046
pm_get_data_mythread, pm_get_tdata_mythread or pm_get_Tdata_mythread Subroutine . . . . .	1047

pm_get_data_mythread_mx or pm_get_tdata_mythread_mx Subroutine . . . . .	1049
pm_get_data_pgroup, pm_get_tdata_pgroup and pm_get_Tdata_pgroup Subroutine . . . . .	1050
pm_get_data_pgroup_mx and pm_get_tdata_pgroup_mx Subroutine . . . . .	1052
pm_get_data_pthread, pm_get_tdata_pthread or pm_get_Tdata_pthread Subroutine . . . . .	1054
pm_get_data_pthread_mx or pm_get_tdata_pthread_mx Subroutine . . . . .	1055
pm_get_data_thread, pm_get_tdata_thread or pm_get_Tdata_thread Subroutine . . . . .	1057
pm_get_data_thread_mx or pm_get_tdata_thread_mx Subroutine . . . . .	1059
pm_get_proctype Subroutine . . . . .	1060
pm_get_program Subroutine . . . . .	1061
pm_get_program_group Subroutine . . . . .	1062
pm_get_program_group_mx and pm_get_program_group_mm Subroutines . . . . .	1064
pm_get_program_mx and pm_get_program_mm Subroutines . . . . .	1066
pm_get_program_mygroup Subroutine . . . . .	1068
pm_get_program_mygroup_mx and pm_get_program_mygroup_mm Subroutines . . . . .	1069
pm_get_program_mythread Subroutine . . . . .	1071
pm_get_program_mythread_mx and pm_get_program_mythread_mm Subroutines . . . . .	1072
pm_get_program_pgroup Subroutine . . . . .	1074
pm_get_program_pgroup_mx and pm_get_program_pgroup_mm Subroutines . . . . .	1075
pm_get_program_pthread Subroutine . . . . .	1077
pm_get_program_pthread_mx and pm_get_program_pthread_mm Subroutines . . . . .	1078
pm_get_program_thread Subroutine . . . . .	1080
pm_get_program_thread_mx and pm_get_program_thread_mm Subroutines . . . . .	1082
pm_init Subroutine . . . . .	1084
pm_initialize Subroutine . . . . .	1085
pm_reset_data Subroutine . . . . .	1087
pm_reset_data_group Subroutine . . . . .	1088
pm_reset_data_mygroup Subroutine . . . . .	1090
pm_reset_data_mythread Subroutine . . . . .	1090
pm_reset_data_pgroup Subroutine . . . . .	1091
pm_reset_data_pthread Subroutine . . . . .	1093
pm_reset_data_thread Subroutine . . . . .	1094
pm_set_program Subroutine . . . . .	1095
pm_set_program_group Subroutine . . . . .	1096
pm_set_program_group_mx and pm_set_program_group_mm Subroutines . . . . .	1098
pm_set_program_mx and pm_set_program_mm Subroutines . . . . .	1100
pm_set_program_mygroup Subroutine . . . . .	1102
pm_set_program_mygroup_mx and pm_set_program_mygroup_mm Subroutines . . . . .	1103
pm_set_program_mythread Subroutine . . . . .	1106
pm_set_program_mythread_mx and pm_set_program_mythread_mm Subroutines . . . . .	1107
pm_set_program_pgroup Subroutine . . . . .	1109
pm_set_program_pgroup_mx and pm_set_program_pgroup_mm Subroutines . . . . .	1111
pm_set_program_pthread Subroutine . . . . .	1113
pm_set_program_pthread_mx and pm_set_program_pthread_mm Subroutines . . . . .	1115
pm_set_program_thread Subroutine . . . . .	1117
pm_set_program_thread_mx and pm_set_program_thread_mm Subroutines . . . . .	1119
pm_start and pm_tstart Subroutine . . . . .	1121
pm_start_group and pm_tstart_group Subroutine . . . . .	1122
pm_start_mygroup and pm_tstart_mygroup Subroutine . . . . .	1123
pm_start_mythread and pm_tstart_mythread Subroutine . . . . .	1125
pm_start_pgroup and pm_tstart_pgroup Subroutine . . . . .	1126
pm_start_pthread and pm_tstart_pthread Subroutine . . . . .	1127
pm_start_thread and pm_tstart_thread Subroutine . . . . .	1129
pm_stop and pm_tstop Subroutine . . . . .	1130
pm_stop_group and pm_tstop_group Subroutine . . . . .	1131
pm_stop_mygroup and pm_tstop_mygroup Subroutine . . . . .	1132
pm_stop_mythread and pm_tstop_mythread Subroutine . . . . .	1133

pm_stop_pgroup and pm_tstop_pgroup Subroutine . . . . .	1134
pm_stop_pthread and pm_tstop_pthread Subroutine . . . . .	1136
pm_stop_thread and pm_tstop_thread Subroutine . . . . .	1137
poll Subroutine . . . . .	1138
pollset_create, pollset_ctl, pollset_destroy, pollset_poll, and pollset_query Subroutines . . . . .	1141
popen Subroutine . . . . .	1144
posix_fadvise Subroutine . . . . .	1145
posix_fallocate Subroutine . . . . .	1146
posix_madvise Subroutine . . . . .	1147
posix_openpt Subroutine . . . . .	1148
posix_spawn or posix_spawnnp Subroutine . . . . .	1149
posix_spawn_file_actions_addclose or posix_spawn_file_actions_addopen Subroutine . . . . .	1153
posix_spawn_file_actions_adddup2 Subroutine. . . . .	1154
posix_spawn_file_actions_destroy or posix_spawn_file_actions_init Subroutine . . . . .	1155
posix_spawnattr_destroy or posix_spawnattr_init Subroutine . . . . .	1156
posix_spawnattr_getflags or posix_spawnattr_setflags Subroutine . . . . .	1157
posix_spawnattr_getpgroup or posix_spawnattr_setpgroup Subroutine . . . . .	1158
posix_spawnattr_getschedparam or posix_spawnattr_setschedparam Subroutine . . . . .	1159
posix_spawnattr_getschedpolicy or posix_spawnattr_setschedpolicy Subroutine . . . . .	1160
posix_spawnattr_getsigdefault or posix_spawnattr_setsigdefault Subroutine . . . . .	1161
posix_spawnattr_getsigmask or posix_spawnattr_setsigmask Subroutine . . . . .	1162
posix_trace_getnext_event, posix_trace_timedgetnext_event, posix_trace_trygetnext_event Subroutine . . . . .	1163
powf, powl, or pow Subroutine . . . . .	1166
printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine . . . . .	1168
profil Subroutine . . . . .	1176
proj_execve Subroutine . . . . .	1178
projdballoc Subroutine . . . . .	1179
projdbfinit Subroutine . . . . .	1179
projdbfree Subroutine . . . . .	1181
psdanger Subroutine . . . . .	1181
psignal Subroutine or sys_siglist Vector . . . . .	1182
pthdb_attr, pthdb_cond, pthdb_condattr, pthdb_key, pthdb_mutex, pthdb_mutexattr, pthdb_pthread, pthdb_pthread_key, pthdb_rwlock, or pthdb_rwlockattr Subroutine . . . . .	1183
pthdb_attr_detachstate, pthdb_attr_addr, pthdb_attr_guardsize, pthdb_attr_inheritsched, pthdb_attr_schedparam, pthdb_attr_schedpolicy, pthdb_attr_schedpriority, pthdb_attr_scope, pthdb_attr_stackaddr, pthdb_attr_stacksize, or pthdb_attr_suspendstate Subroutine. . . . .	1185
pthdb_condattr_pshared, or pthdb_condattr_addr Subroutine . . . . .	1187
pthdb_cond_addr, pthdb_cond_mutex or pthdb_cond_pshared Subroutine. . . . .	1188
pthdb_mutexattr_addr, pthdb_mutexattr_prioceiling, pthdb_mutexattr_protocol, pthdb_mutexattr_pshared or pthdb_mutexattr_type Subroutine . . . . .	1189
pthdb_mutex_addr, pthdb_mutex_lock_count, pthdb_mutex_owner, pthdb_mutex_pshared, pthdb_mutex_prioceiling, pthdb_mutex_protocol, pthdb_mutex_state or pthdb_mutex_type Subroutine . . . . .	1191
pthdb_mutex_waiter, pthdb_cond_waiter, pthdb_rwlock_read_waiter or pthdb_rwlock_write_waiter Subroutine . . . . .	1193
pthdb_pthread_arg Subroutine . . . . .	1194
pthdb_pthread_context or pthdb_pthread_setcontext Subroutine . . . . .	1197
pthdb_pthread_hold, pthdb_pthread_holdstate or pthdb_pthread_unhold Subroutine . . . . .	1198
pthdb_pthread_sigmask, pthdb_pthread_sigpend or pthdb_pthread_sigwait Subroutine . . . . .	1199
pthdb_pthread_specific Subroutine . . . . .	1200
pthdb_pthread_tid or pthdb_tid_pthread Subroutine . . . . .	1201
pthdb_rwlockattr_addr, or pthdb_rwlockattr_pshared Subroutine . . . . .	1202
pthdb_rwlock_addr, pthdb_rwlock_lock_count, pthdb_rwlock_owner, pthdb_rwlock_pshared or pthdb_rwlock_state Subroutine. . . . .	1203
pthdb_session_committed Subroutines. . . . .	1205

pthread_atfork Subroutine . . . . .	1208
pthread_attr_destroy Subroutine . . . . .	1209
pthread_attr_getguardsize or pthread_attr_setguardsize Subroutines. . . . .	1210
pthread_attr_getinheritsched, pthread_attr_setinheritsched Subroutine . . . . .	1212
pthread_attr_getschedparam Subroutine . . . . .	1213
pthread_attr_getschedpolicy, pthread_attr_setschedpolicy Subroutine . . . . .	1214
pthread_attr_getstackaddr Subroutine . . . . .	1215
pthread_attr_getstacksize Subroutine . . . . .	1216
pthread_attr_init Subroutine . . . . .	1217
pthread_attr_getdetachstate or pthread_attr_setdetachstate Subroutines . . . . .	1218
pthread_attr_getscope and pthread_attr_setscope Subroutines . . . . .	1219
pthread_attr_getukeyset_np or pthread_attr_setukeyset_np Subroutine . . . . .	1221
pthread_attr_getschedparam Subroutine . . . . .	1222
pthread_attr_setstackaddr Subroutine . . . . .	1223
pthread_attr_setstacksize Subroutine . . . . .	1224
pthread_attr_setsuspendstate_np and pthread_attr_getsuspendstate_np Subroutine . . . . .	1225
pthread_barrier_destroy or pthread_barrier_init Subroutine . . . . .	1226
pthread_barrier_wait Subroutine . . . . .	1228
pthread_barrierattr_destroy or pthread_barrierattr_init Subroutine . . . . .	1229
pthread_barrierattr_getpshared or pthread_barrierattr_setpshared Subroutine . . . . .	1230
pthread_cancel Subroutine . . . . .	1231
pthread_cleanup_pop or pthread_cleanup_push Subroutine . . . . .	1232
pthread_cond_destroy or pthread_cond_init Subroutine . . . . .	1233
PTHREAD_COND_INITIALIZER Macro . . . . .	1235
pthread_cond_signal or pthread_cond_broadcast Subroutine . . . . .	1235
pthread_cond_wait or pthread_cond_timedwait Subroutine . . . . .	1236
pthread_condattr_destroy or pthread_condattr_init Subroutine . . . . .	1238
pthread_condattr_getclock, pthread_condattr_setclock Subroutine. . . . .	1240
pthread_condattr_getpshared Subroutine . . . . .	1241
pthread_condattr_setpshared Subroutine . . . . .	1242
pthread_create Subroutine . . . . .	1243
pthread_create_withcred_np Subroutine . . . . .	1245
pthread_delay_np Subroutine . . . . .	1246
pthread_equal Subroutine . . . . .	1247
pthread_exit Subroutine . . . . .	1248
pthread_get_expiration_np Subroutine . . . . .	1250
pthread_getconcurrency or pthread_setconcurrency Subroutine . . . . .	1251
pthread_getcpuclockid Subroutine . . . . .	1252
pthread_getrusage_np Subroutine . . . . .	1252
pthread_getschedparam Subroutine. . . . .	1255
pthread_getspecific or pthread_setspecific Subroutine . . . . .	1256
pthread_getthrs_np Subroutine . . . . .	1258
pthread_getunique_np Subroutine . . . . .	1261
pthread_join or pthread_detach Subroutine . . . . .	1262
pthread_key_create Subroutine . . . . .	1263
pthread_key_delete Subroutine . . . . .	1264
pthread_kill Subroutine . . . . .	1265
pthread_lock_global_np Subroutine . . . . .	1266
pthread_mutex_init or pthread_mutex_destroy Subroutine. . . . .	1267
pthread_mutex_getprioceiling or pthread_mutex_setprioceiling Subroutine. . . . .	1269
PTHREAD_MUTEX_INITIALIZER Macro . . . . .	1270
pthread_mutex_lock, pthread_mutex_trylock, or pthread_mutex_unlock Subroutine . . . . .	1270
pthread_mutex_timedlock Subroutine . . . . .	1272
pthread_mutexattr_destroy or pthread_mutexattr_init Subroutine . . . . .	1273
pthread_mutexattr_getkind_np Subroutine . . . . .	1275
pthread_mutexattr_getprioceiling or pthread_mutexattr_setprioceiling Subroutine . . . . .	1276

pthread_mutexattr_getprotocol or pthread_mutexattr_setprotocol Subroutine . . . . .	1277
pthread_mutexattr_getpshared or pthread_mutexattr_setpshared Subroutine . . . . .	1279
pthread_mutexattr_gettype or pthread_mutexattr_settype Subroutine . . . . .	1280
pthread_mutexattr_setkind_np Subroutine . . . . .	1281
pthread_once Subroutine. . . . .	1283
PTHREAD_ONCE_INIT Macro . . . . .	1284
pthread_rwlock_init or pthread_rwlock_destroy Subroutine . . . . .	1284
pthread_rwlock_rdlock or pthread_rwlock_tryrdlock Subroutines . . . . .	1286
pthread_rwlock_timedrdlock Subroutine . . . . .	1287
pthread_rwlock_timedwrlock Subroutine . . . . .	1289
pthread_rwlock_unlock Subroutine . . . . .	1290
pthread_rwlock_wrlock or pthread_rwlock_trywrlock Subroutines . . . . .	1291
pthread_rwlockattr_init or pthread_rwlockattr_destroy Subroutines. . . . .	1293
pthread_rwlockattr_getpshared or pthread_rwlockattr_setpshared Subroutines . . . . .	1294
pthread_self Subroutine . . . . .	1295
pthread_setcancelstate, pthread_setcanceltype, or pthread_testcancel Subroutines . . . . .	1296
pthread_setschedparam Subroutine . . . . .	1297
pthread_setschedprio Subroutine . . . . .	1299
pthread_sigmask Subroutine . . . . .	1300
pthread_signal_to_cancel_np Subroutine . . . . .	1300
pthread_spin_destroy or pthread_spin_init Subroutine . . . . .	1301
pthread_spin_lock or pthread_spin_trylock Subroutine . . . . .	1302
pthread_spin_unlock Subroutine . . . . .	1303
pthread_suspend_np, pthread_unsuspend_np and pthread_continue_np Subroutine . . . . .	1304
pthread_unlock_global_np Subroutine . . . . .	1305
pthread_yield Subroutine . . . . .	1306
ptrace, ptracex, ptrace64 Subroutine . . . . .	1307
ptsname Subroutine . . . . .	1320
putc, putchar, fputc, or putw Subroutine . . . . .	1321
putconfattrs Subroutine . . . . .	1323
putenv Subroutine . . . . .	1326
putgrent Subroutine. . . . .	1326
putgroupattrs Subroutine . . . . .	1327
puts or fputs Subroutine . . . . .	1331
putuserattrs Subroutine . . . . .	1333
putuserpwx Subroutine . . . . .	1337
putwc, putwchar, or fputwc Subroutine . . . . .	1339
putws or fputws Subroutine . . . . .	1341
pwdrestrict_method Subroutine . . . . .	1342
<b>Appendix A. Base Operating System Error Codes for Services That Require Path-Name</b>	
<b>Resolution</b> . . . . .	1345
Related Information. . . . .	1345
<b>Appendix B. ODM Error Codes</b> . . . . .	1347
Related Information. . . . .	1348
<b>Appendix C. Notices</b> . . . . .	1349
Trademarks. . . . .	1350
<b>Index</b> . . . . .	1351



---

## About This Book

This book provides experienced C programmers with complete detailed information about Base Operating System runtime services for the AIX operating system. Runtime services are listed alphabetically, and complete descriptions are given for them. This volume contains AIX services that begin with the letters A through P. To use the book effectively, you should be familiar with commands, system calls, subroutines, file formats, and special files. This publication is also available on the documentation CD that is shipped with the operating system.

This book is part of the six-volume technical reference set, *AIX 5L Version 5.3 Technical Reference*, that provides information on system calls, kernel extension calls, and subroutines in the following volumes:

- *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 1* and *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2* provide information on system calls, subroutines, functions, macros, and statements associated with base operating system runtime services.
- *AIX 5L Version 5.3 Technical Reference: Communications Volume 1* and *AIX 5L Version 5.3 Technical Reference: Communications Volume 2* provide information on entry points, functions, system calls, subroutines, and operations related to communications services.
- *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 1* and *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 2* provide information about kernel services, device driver operations, file system operations, subroutines, the configuration subsystem, the communications subsystem, the low function terminal (LFT) subsystem, the logical volume subsystem, the M-audio capture and playback adapter subsystem, the printer subsystem, the SCSI subsystem, and the serial DASD subsystem.

---

## Highlighting

The following highlighting conventions are used in this book:

<b>Bold</b>	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
<i>Italics</i>	Identifies parameters whose actual names or values are to be supplied by the user.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

---

## Case-Sensitivity in AIX

Everything in the AIX operating system is case-sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is "not found." Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

---

## ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

---

## 32-Bit and 64-Bit Support for the Single UNIX Specification

Beginning with Version 5.2, the operating system is designed to support The Open Group's Single UNIX Specification Version 3 (UNIX 03) for portability of UNIX-based operating systems. Many new interfaces, and some current ones, have been added or enhanced to meet this specification, making Version 5.2 even more open and portable for applications, while remaining compatible with previous releases of AIX. To determine the proper way to develop a UNIX 03-portable application, you may need to refer to The Open Group's UNIX 03 specification, which can be accessed online or downloaded from <http://www.unix.org/> .

---

## Related Publications

The following books contain information about or related to application programming interfaces:

- *Operating system and device management*
- *Networks and communication management*
- *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*
- *AIX 5L Version 5.3 Communications Programming Concepts*
- *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*
- *AIX 5L Version 5.3 Files Reference*



---

# Base Operating System (BOS) Runtime Services (A-P)

---

## a64I or I64a Subroutine

### Purpose

Converts between long integers and base-64 ASCII strings.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
long a64I ( String)  
char *String;
```

```
char *I64a ( LongInteger )  
long LongInteger;
```

### Description

The **a64I** and **I64a** subroutines maintain numbers stored in base-64 ASCII characters. This is a notation in which long integers are represented by up to 6 characters, each character representing a digit in a base-64 notation.

The following characters are used to represent digits:

Character	Description
.	Represents 0.
/	Represents 1.
0 -9	Represents the numbers 2-11.
A-Z	Represents the numbers 12-37.
a-z	Represents the numbers 38-63.

### Parameters

*String* Specifies the address of a null-terminated character string.  
*LongInteger* Specifies a long value to convert.

### Return Values

The **a64I** subroutine takes a pointer to a null-terminated character string containing a value in base-64 representation and returns the corresponding **long** value. If the string pointed to by the *String* parameter contains more than 6 characters, the **a64I** subroutine uses only the first 6.

Conversely, the **I64a** subroutine takes a **long** parameter and returns a pointer to the corresponding base-64 representation. If the *LongInteger* parameter is a value of 0, the **I64a** subroutine returns a pointer to a null string.

The value returned by the **I64a** subroutine is a pointer into a static buffer, the contents of which are overwritten by each call.

If the *\*String* parameter is a null string, the **a64I** subroutine returns a value of 0L.

If *LongInteger* is 0L, the **I64a** subroutine returns a pointer to a null string.

## Related Information

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

List of Multithread Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## abort Subroutine

### Purpose

Sends a **SIGIOT** signal to end the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
int abort (void)
```

### Description

The **abort** subroutine sends a **SIGIOT** signal to the current process to terminate the process and produce a memory dump. If the signal is caught and the signal handler does not return, the **abort** subroutine does not produce a memory dump.

If the **SIGIOT** signal is neither caught nor ignored, and if the current directory is writable, the system produces a memory dump in the **core** file in the current directory and prints an error message.

The abnormal-termination processing includes the effect of the **fclose** subroutine on all open streams and message-catalog descriptors, and the default actions defined as the **SIGIOT** signal. The **SIGIOT** signal is sent in the same manner as that sent by the **raise** subroutine with the argument **SIGIOT**.

The status made available to the **wait** or **waitpid** subroutine by the **abort** subroutine is the same as a process terminated by the **SIGIOT** signal. The **abort** subroutine overrides blocking or ignoring the **SIGIOT** signal.

**Note:** The **SIGABRT** signal is the same as the **SIGIOT** signal.

### Return Values

The **abort** subroutine does not return a value.

## Related Information

The **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243), **atexit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243), or **\_exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fclose** (“fclose or fflush Subroutine” on page 253) subroutine, **kill** (“kill or killpg Subroutine” on page 580), or **killpg** (“kill or killpg Subroutine” on page 580) subroutine, **raise** subroutine, **sigaction**, **sigvec**, **signal** subroutine, **wait** or **waitpid** subroutine.

The **dbx** command.

## abs, div, labs, ldiv, imul\_dbl, umul\_dbl, llabs, or lldiv Subroutine

### Purpose

Computes absolute value, division, and double precision multiplication of integers.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
int abs ( i )
int i;

#include <stdlib.h>
long labs ( i )
long i;

#include <stdlib.h>

div_t div ( Numerator, Denominator)
int Numerator: Denominator;
#include <stdlib.h>

void imul_dbl ( i, j, Result)
long i, j;
long *Result;
#include <stdlib.h>

ldiv_t ldiv (Numerator, Denominator)
long Numerator: Denominator;
#include <stdlib.h>

void umul_dbl (i, j, Result)
unsigned long i, j;
unsigned long *Result;
#include <stdlib.h>

long long int llabs(i)
long long int i;
#include <stdlib.h>

lldiv_t lldiv (Numerator, Denominator)
long long int Numerator, Denominator;
```

### Description

The **abs** subroutine returns the absolute value of its integer operand.

**Note:** A two's-complement integer can hold a negative number whose absolute value is too large for the integer to hold. When given this largest negative value, the **abs** subroutine returns the same value.

The **div** subroutine computes the quotient and remainder of the division of the number represented by the *Numerator* parameter by that specified by the *Denominator* parameter. If the division is inexact, the sign of the resulting quotient is that of the algebraic quotient, and the magnitude of the resulting quotient is the largest integer less than the magnitude of the algebraic quotient. If the result cannot be represented (for example, if the denominator is 0), the behavior is undefined.

The **labs** and **ldiv** subroutines are included for compatibility with the ANSI C library, and accept long integers as parameters, rather than as integers.

The **imul\_dbl** subroutine computes the product of two signed longs, *i* and *j*, and stores the double long product into an array of two signed longs pointed to by the *Result* parameter.

The **umul\_dbl** subroutine computes the product of two unsigned longs, *i* and *j*, and stores the double unsigned long product into an array of two unsigned longs pointed to by the *Result* parameter.

The **llabs** and **lldiv** subroutines compute the absolute value and division of long long integers. These subroutines operate under the same restrictions as the **abs** and **div** subroutines.

**Note:** When given the largest negative value, the **llabs** subroutine (like the **abs** subroutine) returns the same value.

## Parameters

<i>i</i>	Specifies, for the <b>abs</b> subroutine, some integer; for <b>labs</b> and <b>imul_dbl</b> , some long integer; for the <b>umul_dbl</b> subroutine, some unsigned long integer; for the <b>llabs</b> subroutine, some long long integer.
<i>Numerator</i>	Specifies, for the <b>div</b> subroutine, some integer; for the <b>ldiv</b> subroutine, some long integer; for <b>lldiv</b> , some long long integer.
<i>j</i>	Specifies, for the <b>imul_dbl</b> subroutine, some long integer; for the <b>umul_dbl</b> subroutine, some unsigned long integer.
<i>Denominator</i>	Specifies, for the <b>div</b> subroutine, some integer; for the <b>ldiv</b> subroutine, some long integer; for <b>lldiv</b> , some long long integer.
<i>Result</i>	Specifies, for the <b>imul_dbl</b> subroutine, some long integer; for the <b>umul_dbl</b> subroutine, some unsigned long integer.

## Return Values

The **abs**, **labs**, and **llabs** subroutines return the absolute value. The **imul\_dbl** and **umul\_dbl** subroutines have no return values. The **div** subroutine returns a structure of type **div\_t**. The **ldiv** subroutine returns a structure of type **ldiv\_t**, comprising the quotient and the remainder. The structure is displayed as:

```
struct ldiv_t {
    int quot; /* quotient */
    int rem; /* remainder */
};
```

The **lldiv** subroutine returns a structure of type **lldiv\_t**, comprising the quotient and the remainder.

---

## access, accessx, or faccessx Subroutine

### Purpose

Determines the accessibility of a file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int access ( PathName, Mode)
char *PathName;
```

```
int Mode;
```

```
int accessx (PathName, Mode, Who)
```

```
char *PathName;
```

```
int Mode, Who;
```

```
int faccessx ( FileDescriptor, Mode, Who)
```

```
int FileDescriptor;
```

```
int Mode, Who;
```

## Description

The **access**, **accessx**, and **faccessx** subroutines determine the accessibility of a file system object. The **accessx** and **faccessx** subroutines allow the specification of a class of users or processes for whom access is to be checked.

The caller must have search permission for all components of the *PathName* parameter.

## Parameters

*PathName*

Specifies the path name of the file. If the *PathName* parameter refers to a symbolic link, the **access** subroutine returns information about the file pointed to by the symbolic link.

*FileDescriptor*

Specifies the file descriptor of an open file.

*Mode*

Specifies the access modes to be checked. This parameter is a bit mask containing 0 or more of the following values, which are defined in the **sys/access.h** file:

**R\_OK** Check read permission.

**W\_OK** Check write permission.

**X\_OK** Check execute or search permission.

**F\_OK** Check the existence of a file.

If none of these values are specified, the existence of a file is checked.

*Who*

Specifies the class of users for whom access is to be checked. This parameter must be one of the following values, which are defined in the **sys/access.h** file:

### **ACC\_SELF**

Determines if access is permitted for the current process. The effective user and group IDs, the concurrent group set and the privilege of the current process are used for the calculation.

### **ACC\_INVOKER**

Determines if access is permitted for the invoker of the current process. The real user and group IDs, the concurrent group set, and the privilege of the invoker are used for the calculation.

**Note:** The expression **access (PathName, Mode)** is equivalent to **accessx (PathName, Mode, ACC\_INVOKER)**.

### **ACC\_OTHERS**

Determines if the specified access is permitted for any user other than the object owner. The *Mode* parameter must contain only one of the valid modes. Privilege is not considered in the calculation.

### **ACC\_ALL**

Determines if the specified access is permitted for all users. The *Mode* parameter must contain only one of the valid modes. Privilege is not considered in the calculation .

**Note:** The **accessx** subroutine shows the same behavior by both the user and root with **ACC\_ALL**.

## Return Values

If the requested access is permitted, the **access**, **accessx**, and **faccessx** subroutines return a value of 0. If the requested access is not permitted or the function call fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

The **access** subroutine indicates success for **X\_OK** even if none of the execute file permission bits are set.

## Error Codes

The **access** and **accessx** subroutines fail if one or more of the following are true:

<b>EACCES</b>	Search permission is denied on a component of the <i>PathName</i> prefix.
<b>EFAULT</b>	The <i>PathName</i> parameter points to a location outside the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>PathName</i> parameter.
<b>ENOENT</b>	A component of the <i>PathName</i> does not exist or the process has the <b>disallow truncation</b> attribute set.
<b>ENOTDIR</b>	A component of the <i>PathName</i> is not a directory.
<b>ESTALE</b>	The process root or current directory is located in a virtual file system that has been unmounted.
<b>ENOENT</b>	The named file does not exist.
<b>ENOENT</b>	The <i>PathName</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENAMETOOLONG</b>	A component of the <i>PathName</i> parameter exceeded 255 characters or the entire <i>PathName</i> parameter exceeded 1022 characters.

The **faccessx** subroutine fails if the following is true:

<b>EBADF</b>	The value of the <i>FileDescriptor</i> parameter is not valid.
--------------	--

The **access**, **accessx**, and **faccessx** subroutines fail if one or more of the following is true:

<b>EIO</b>	An I/O error occurred during the operation.
<b>EACCES</b>	The file protection does not allow the requested access.
<b>EROFS</b>	Write access is requested for a file on a read-only file system.

If Network File System (NFS) is installed on your system, the **accessx** and **faccessx** subroutines can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
<b>ETXTBSY</b>	Write access is requested for a shared text file that is being executed.
<b>EINVAL</b>	The value of the <i>Mode</i> argument is invalid.

## Related Information

The **acl\_get** (“acl\_get or acl\_fget Subroutine” on page 10) subroutine, **chacl** (“chacl or fchacl Subroutine” on page 144) subroutine, **statx** subroutine, **statacl** subroutine.

The **aclget** command, **aclput** command, **chmod** command, **chown** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## acct Subroutine

### Purpose

Enables and disables process accounting.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int acct ( Path)
char *Path;
```

### Description

The **acct** subroutine enables the accounting routine when the *Path* parameter specifies the path name of the file to which an accounting record is written for each process that terminates. When the *Path* parameter is a 0 or null value, the **acct** subroutine disables the accounting routine.

If the *Path* parameter refers to a symbolic link, the **acct** subroutine causes records to be written to the file pointed to by the symbolic link.

If Network File System (NFS) is installed on your system, the accounting file can reside on another node.

**Note:** To ensure accurate accounting, each node must have its own accounting file. Although no two nodes should share accounting files, a node's accounting files can be located on any node in the network.

The calling process must have root user authority to use the **acct** subroutine.

### Parameters

*Path*            Specifies a pointer to the path name of the file or a null pointer.

### Return Values

Upon successful completion, the **acct** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the global variable **errno** is set to indicate the error.

### Error Codes

The **acct** subroutine is unsuccessful if one or more of the following are true:

<b>EACCES</b>	Write permission is denied for the named accounting file.
<b>EACCES</b>	The file named by the <i>Path</i> parameter is not an ordinary file.
<b>EBUSY</b>	An attempt is made to enable accounting when it is already enabled.
<b>ENOENT</b>	The file named by the <i>Path</i> parameter does not exist.
<b>EPERM</b>	The calling process does not have root user authority.
<b>EROFS</b>	The named file resides on a read-only file system.

If NFS is installed on the system, the **acct** subroutine is unsuccessful if the following is true:

**ETIMEDOUT**            The connection timed out.

---

## acl\_chg or acl\_fchg Subroutine

### Purpose

Changes the AIXC ACL type access control information on a file.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/access.h>
```

```
int acl_chg (Path, How, Mode, Who)
```

```
char * Path;
```

```
int How;
```

```
int Mode;
```

```
int Who;
```

```
int acl_fchg (FileDescriptor, How, Mode, Who)
```

```
int FileDescriptor;
```

```
int How;
```

```
int Mode;
```

```
int Who;
```

### Description

The **acl\_chg** and **acl\_fchg** subroutines modify the AIXC ACL-type-based access control information of a specified file. This call can fail for file system objects with any non-AIXC ACL.

### Parameters

*FileDescriptor*

Specifies the file descriptor of an open file.

*How*

Specifies how the permissions are to be altered for the affected entries of the Access Control List (ACL). This parameter takes one of the following values:

**ACC\_PERMIT**

Allows the types of access included in the *Mode* parameter.

**ACC\_DENY**

Denies the types of access included in the *Mode* parameter.

**ACC\_SPECIFY**

Grants the access modes included in the *Mode* parameter and restricts the access modes not included in the *Mode* parameter.

*Mode*

Specifies the access modes to be changed. The *Mode* parameter is a bit mask containing zero or more of the following values:

**R\_ACC**

Allows read permission.

**W\_ACC**

Allows write permission.

**X\_ACC** Allows execute or search permission.

*Path*

Specifies a pointer to the path name of a file.



*Who* Specifies which entries in the ACL are affected. This parameter takes one of the following values:

**ACC\_OBJ\_OWNER**

Changes the owner entry in the base ACL.

**ACC\_OBJ\_GROUP**

Changes the group entry in the base ACL.

**ACC\_OTHERS**

Changes all entries in the ACL except the base entry for the owner.

**ACC\_ALL**

Changes all entries in the ACL.

## Return Values

On successful completion, the **acl\_chg** and **acl\_fchg** subroutines return a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **acl\_chg** subroutine fails and the access control information for a file remains unchanged if one or more of the following is true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

The **acl\_fchg** subroutine fails and the file permissions remain unchanged if the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> value is not valid.
--------------	---

The **acl\_chg** or **acl\_fchg** subroutine fails and the access control information for a file remains unchanged if one or more of the following is true:

<b>EINVAL</b>	The <i>How</i> parameter is not one of <b>ACC_PERMIT</b> , <b>ACC_DENY</b> , or <b>ACC_SPECIFY</b> .
<b>EINVAL</b>	The <i>Who</i> parameter is not <b>ACC_OWNER</b> , <b>ACC_GROUP</b> , <b>ACC_OTHERS</b> , or <b>ACC_ALL</b> .
<b>EROFS</b>	The named file resides on a read-only file system.

The **acl\_chg** or **acl\_fchg** subroutine fails and the access control information for a file remains unchanged if one or more of the following is true:

<b>EIO</b>	An I/O error occurred during the operation.
<b>EPERM</b>	The effective user ID does not match the ID of the owner of the file and the invoker does not have root user authority.

If Network File System (NFS) is installed on your system, the **acl\_chg** and **acl\_fchg** subroutines can also fail if the following is true:

**ETIMEDOUT**                      The connection timed out.

## Related Information

The **acl\_get** (“acl\_get or acl\_fget Subroutine”) subroutine, **acl\_put** (“acl\_put or acl\_fput Subroutine” on page 12) subroutine, **acl\_set** (“acl\_set or acl\_fset Subroutine” on page 14) subroutine, **chacl** (“chacl or fchacl Subroutine” on page 144) subroutine, **chmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## acl\_get or acl\_fget Subroutine

### Purpose

Gets the access control information of a file if the ACL associated is of the AIXC type.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/access.h>
```

```
char *acl_get (Path)
char * Path;
```

```
char *acl_fget (FileDescriptor)
int FileDescriptor;
```

### Description

The **acl\_get** and **acl\_fget** subroutines retrieve the access control information for a file system object. This information is returned in a buffer pointed to by the return value. The structure of the data in this buffer is unspecified. The value returned by these subroutines should be used only as an argument to the **acl\_put** or **acl\_fput** subroutines to copy or restore the access control information. Note that **acl\_get** and **acl\_fget** subroutines could fail if the ACL associated with the file system object is of a different type than AIXC. It is recommended that applications make use of **aclx\_get** and **aclx\_fget** subroutines to retrieve the ACL.

The buffer returned by the **acl\_get** and **acl\_fget** subroutines is in allocated memory. After usage, the caller should deallocate the buffer using the **free** subroutine.

### Parameters

<i>Path</i>	Specifies the path name of the file.
<i>FileDescriptor</i>	Specifies the file descriptor of an open file.

## Return Values

On successful completion, the **acl\_get** and **acl\_fget** subroutines return a pointer to the buffer containing the access control information. Otherwise, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **acl\_get** subroutine fails if one or more of the following are true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or the process has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

The **acl\_fget** subroutine fails if the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor.
--------------	---

The **acl\_get** or **acl\_fget** subroutine fails if the following is true:

<b>EIO</b>	An I/O error occurred during the operation.
------------	---

If Network File System (NFS) is installed on your system, the **acl\_get** and **acl\_fget** subroutines can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

<b>Access Control</b>	The invoker must have search permission for all components of the <i>Path</i> prefix.
<b>Audit Events</b>	None.

## Related Information

The **acl\_chg** or **acl\_fchg** (“**acl\_chg** or **acl\_fchg** Subroutine” on page 8) subroutine, **acl\_put** or **acl\_fput** (“**acl\_put** or **acl\_fput** Subroutine” on page 12) subroutine, **acl\_set** or **acl\_fset** (“**acl\_set** or **acl\_fset** Subroutine” on page 14) subroutine, **chacl** (“**chacl** or **fchacl** Subroutine” on page 144) subroutine, **chmod** (“**chmod** or **fchmod** Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine.

“**aclx\_get** or **aclx\_fget** Subroutine” on page 17, “**aclx\_put** or **aclx\_fput** Subroutine” on page 25.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## acl\_put or acl\_fput Subroutine

### Purpose

Sets AIXC ACL type access control information of a file.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/access.h>
```

```
int acl_put (Path, Access, Free)
```

```
char * Path;
```

```
char * Access;
```

```
int Free;
```

```
int acl_fput (FileDescriptor, Access, Free)
```

```
int FileDescriptor;
```

```
char * Access;
```

```
int Free;
```

### Description

The **acl\_put** and **acl\_fput** subroutines set the access control information of a file system object. This information is contained in a buffer returned by a call to the **acl\_get** or **acl\_fget** subroutine. The structure of the data in this buffer is unspecified. However, the entire Access Control List (ACL) for a file cannot exceed one memory page (4096 bytes) in size. Note that **acl\_put/acl\_fput** operation could fail if the existing ACL associated with the file system object is of a different kind or if the underlying physical file system does not support AIXC ACL type. It is recommended that applications make use of **aclx\_put** and **aclx\_fput** subroutines to set the ACL instead of **acl\_put/acl\_fput** routines.

### Parameters

<i>Path</i>	Specifies the path name of a file.
<i>FileDescriptor</i>	Specifies the file descriptor of an open file.
<i>Access</i>	Specifies a pointer to the buffer containing the access control information.
<i>Free</i>	Specifies whether the buffer space is to be deallocated. The following values are valid:
0	Space is not deallocated.
1	Space is deallocated.

### Return Values

On successful completion, the **acl\_put** and **acl\_fput** subroutines return a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **acl\_put** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>EACCESS</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.

<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

The **acl\_fput** subroutine fails and the file permissions remain unchanged if the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor.
--------------	---

The **acl\_put** or **acl\_fput** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>EINVAL</b>	The <i>Access</i> parameter does not point to a valid access control buffer.
<b>EINVAL</b>	The <i>Free</i> parameter is not 0 or 1.
<b>EIO</b>	An I/O error occurred during the operation.
<b>EROFS</b>	The named file resides on a read-only file system.

If Network File System (NFS) is installed on your system, the **acl\_put** and **acl\_fput** subroutines can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

Auditing Events:

Event	Information
<b>chacl</b>	<i>Path</i>
<b>fchacl</b>	<i>FileDescriptor</i>

## Related Information

The **acl\_chg** (“acl\_chg or acl\_fchg Subroutine” on page 8) subroutine, **acl\_get** (“acl\_get or acl\_fget Subroutine” on page 10) subroutine, **acl\_set** (“acl\_set or acl\_fset Subroutine” on page 14) subroutine, **chacl** (“chacl or fchacl Subroutine” on page 144) subroutine, **chmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine.

“aclx\_get or aclx\_fget Subroutine” on page 17, “aclx\_put or aclx\_fput Subroutine” on page 25.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## acl\_set or acl\_fset Subroutine

### Purpose

Sets the AIXC ACL type access control information of a file.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/access.h>
```

```
int acl_set (Path, OwnerMode, GroupMode, DefaultMode)
char * Path;
int OwnerMode;
int GroupMode;
int DefaultMode;
```

```
int acl_fset (FileDescriptor, OwnerMode, GroupMode, DefaultMode)
int * FileDescriptor;
int OwnerMode;
int GroupMode;
int DefaultMode;
```

### Description

The **acl\_set** and **acl\_fset** subroutines set the base entries of the Access Control List (ACL) of the file. All other entries are discarded. Other access control attributes are left unchanged. Note that if the file system object is associated with any other ACL type access control information, it will be replaced with just the Base mode bits information. It is strongly recommended that applications stop using these interfaces and instead make use of **aclx\_put** and **aclx\_fput** subroutines to set the ACL.

### Parameters

<i>DefaultMode</i>	Specifies the access permissions for the default class.
<i>FileDescriptor</i>	Specifies the file descriptor of an open file.
<i>GroupMode</i>	Specifies the access permissions for the group of the file.
<i>OwnerMode</i>	Specifies the access permissions for the owner of the file.
<i>Path</i>	Specifies a pointer to the path name of a file.

The mode parameters specify the access permissions in a bit mask containing zero or more of the following values:

<b>R_ACC</b>	Authorize read permission.
<b>W_ACC</b>	Authorize write permission.
<b>X_ACC</b>	Authorize execute or search permission.

### Return Values

Upon successful completion, the **acl\_set** and **acl\_fset** subroutines return the value 0. Otherwise, the value -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **acl\_set** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENODIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

The **acl\_fset** subroutine fails and the file permissions remain unchanged if the following is true:

<b>EBADF</b>	The file descriptor <i>FileDescriptor</i> is not valid.
--------------	---

The **acl\_set** or **acl\_fset** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>EIO</b>	An I/O error occurred during the operation.
<b>EPERM</b>	The effective user ID does not match the ID of the owner of the file and the invoker does not have root user authority.
<b>EROFS</b>	The named file resides on a read-only file system.

If Network File System (NFS) is installed on your system, the **acl\_set** and **acl\_fset** subroutines can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

Auditing Events:

Event	Information
<b>chacl</b>	<i>Path</i>
<b>fchacl</b>	<i>FileDescriptor</i>

## Related Information

The **acl\_chg** (“acl\_chg or acl\_fchg Subroutine” on page 8) subroutine, **acl\_get** (“acl\_get or acl\_fget Subroutine” on page 10) subroutine, **acl\_put** (“acl\_put or acl\_fput Subroutine” on page 12) subroutine, **chacl** (“chacl or fchacl Subroutine” on page 144) subroutine, **chmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine.

“aclx\_get or aclx\_fget Subroutine” on page 17, “aclx\_put or aclx\_fput Subroutine” on page 25.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## aclx\_convert Subroutine

### Purpose

Converts the access control information from one ACL type to another.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_convert (from_acl, from_sz, from_type, to_acl, to_sz, to_type, fs_obj_path)
void * from_acl;
size_t from_sz;
acl_type_t from_type;
void * to_acl;
size_t * to_sz;
acl_type_t to_type;
char * fs_obj_path;
```

### Description

The **aclx\_convert** subroutine converts the access control information from the binary input given in *from\_acl* of the ACL type *from\_type* into a binary ACL of the type *to\_type* and stores it in *to\_acl*. Values *from\_type* and *to\_type* can be any ACL types supported in the system.

The ACL conversion takes place with the help of an ACL type-specific algorithm. Because the conversion is approximate, it can result in a potential loss of access control. Therefore, the user of this call must make sure that the converted ACL satisfies the required access controls. The user can manually review the access control information after the conversion for the file system object to ensure that the conversion was successful and satisfied the requirements of the intended access control.

### Parameters

<i>from_acl</i>	Points to the ACL that has to be converted.
<i>from_sz</i>	Indicates the size of the ACL information pointed to by <i>from_acl</i> .
<i>from_type</i>	Indicates the ACL type information of the ACL. The <i>acl_type</i> is 64 bits in size and is unique on the system. If the given <i>acl_type</i> is not supported in the system, this function fails and <b>errno</b> is set to <b>EINVAL</b> .
<i>to_acl</i>	Points to a buffer in which the target binary ACL has to be stored. The amount of memory available in this buffer is indicated by the <i>to_sz</i> parameter.
<i>to_sz</i>	Indicates the amount of memory, in bytes, available in <i>to_acl</i> . If <i>to_sz</i> contains less than the required amount of memory for storing the converted ACL, <i>*to_sz</i> is set to the required amount of memory and <b>ENOSPC</b> is returned by <b>errno</b> .
<i>to_type</i>	Indicates the ACL type to which conversion needs to be done. The ACL type is 64 bits in size and is unique on the system. If the given <i>acl_type</i> is not supported in the system, this function fails and <b>errno</b> is set to <b>EINVAL</b> .
<i>fs_obj_path</i>	File System Object Path for which the ACL conversion is being requested. Gets information about the object, such as whether it is file or directory.



## Return Values

On successful completion, the **aclx\_convert** subroutine returns a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **aclx\_convert** subroutine fails if one or more of the following is true:

<b>EINVAL</b>	Invalid input parameter. The same error can be returned if an invalid <i>acl_type</i> is specified as input to this routine, either in <i>from_type</i> or in <i>to_type</i> . This <b>errno</b> could also be returned if the binary ACL given in <i>from_acl</i> is not the type specified by <i>from_type</i> .
<b>ENOSPC</b>	Insufficient storage space is available in <i>to_acl</i> .

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

Auditing Events: If the auditing subsystem has been properly configured and is enabled, the **aclx\_convert** subroutine generates the following audit record (event) every time the command is executed:

Event	Information
<b>FILE_Acl</b>	Lists access controls.

## Related Information

The **aclget** command, **aclput** command, **aclconvert** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## aclx\_get or aclx\_fget Subroutine

### Purpose

Gets the access control information for a file system object.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_get (Path, ctl_flags, acl_type, acl, acl_sz, mode_info)
char * Path;
uint64_t  ctl_flags;
acl_type_t * acl_type;
void * acl;
size_t * acl_sz;
mode_t * mode_info;
```

```
int aclx_fget (FileDescriptor, ctl_flags, acl_type, acl, acl_sz, mode_info)
int FileDescriptor;
uint64_t  ctl_flags;
acl_type_t * acl_type;
```

```
void * acl;
size_t * acl_sz;
mode_t * mode_info;
```

## Description

The **aclx\_get** and **aclx\_fget** subroutines retrieve the access control information for a file system object in the native ACL format. Native ACL format is the format as defined for the particular ACL type in the system. These subroutines are advanced versions of the **acl\_get** and **acl\_fget** subroutines and should be used instead of the older versions. The **aclx\_get** and **aclx\_fget** subroutines provide for more control for the user to interact with the underlying file system directly.

In the earlier versions (**acl\_get** or **acl\_fget**), OS libraries found out the ACL size from the file system and allocated the required memory buffer space to hold the ACL information. The caller does all this now with the **aclx\_get** and **aclx\_fget** subroutines. Callers are responsible for finding out the size and allocating memory for the ACL information, and later freeing the same memory after it is used. These subroutines allow for an *acl\_type* input and output argument. The data specified in this argument can be set to a particular ACL type and a request for the ACL on the file system object of the same type. Some physical file systems might do emulation to return the ACL type requested, if the ACL type that exists on the file system object is different. If the *acl\_type* pointer points to a data area with a value of **ACL\_ANY** or 0, then the underlying physical file system has to return the type of the ACL associated with the file system object.

The *ctl\_flags* parameter is a bit mask that allows for control over the **aclx\_get** requests.

The value returned by these subroutines can be use as an argument to the **aclx\_get** or **aclx\_fget** subroutines to copy or restore the access control information.

## Parameters

<i>Path</i>	Specifies the path name of the file system object.
<i>FileDescriptor</i>	Specifies the file descriptor of an open file.
<i>ctl_flags</i>	This 64-bit sized bit mask provides control over the ACL retrieval. The following flag values are defined:  <b>GET_ACLINFO_ONLY</b> Gets only the ACL type and length information from the underlying file system. When this bit is set, arguments such as <i>acl</i> and <i>mode_info</i> can be set to NULL. In all other cases, these should be valid buffer pointers (or else an error is returned). If this bit is not specified, then all the other information about the ACL, such as ACL data and mode information, is returned.
<i>acl_type</i>	Points to a buffer that will hold ACL type information. The ACL type is 64 bits in size and is unique on the system. The caller can provide an ACL type in this area and a request for the ACL on the file system object of the same type. If the ACL type requested does not match the one on the file system object, the physical file system might return an error or emulate and provide the ACL information in the ACL type format requested. If the caller does not know the ACL type and wants to retrieve the ACL associated with the file system object, then the caller should set the buffer value pointed to by <i>acl_type</i> to <b>ACL_ANY</b> or 0.
<i>acl</i>	Points to a buffer where the ACL retrieved is stored. The size of this buffer is indicated by the <i>acl_sz</i> parameter.
<i>acl_sz</i>	Indicates the size of the buffer area passed through the <i>acl</i> parameter.
<i>mode_info</i>	Pointer to a buffer where the mode word associated with the file system object is returned. Note that this mode word's meaning and formations depend entirely on the ACL type concerned.

## Return Values

On successful completion, the **aclx\_put** and **aclx\_fput** subroutines return a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **aclx\_get** subroutine fails if one or more of the following is true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

The **aclx\_fget** subroutine fails if the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor.
--------------	---

The **aclx\_get** or **aclx\_fget** subroutine fails if one or more of the following is true:

<b>EINVAL</b>	Invalid input parameter. The same error can be returned if an invalid <i>acl_type</i> is specified as input to this routine.
<b>EIO</b>	An I/O error occurred during the operation.
<b>ENOSPC</b>	Input buffer size <i>acl_sz</i> is not sufficient to store the ACL data in <i>acl</i> .

If Network File System (NFS) is installed on your system, the **aclx\_get** and **aclx\_fget** subroutines can also fail if the following condition is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

Auditing Events: None

## Related Information

The **acl\_chg** (“acl\_chg or acl\_fchg Subroutine” on page 8) subroutine, **acl\_put** (“acl\_get or acl\_fget Subroutine” on page 10) subroutine, **acl\_set** (“acl\_set or acl\_fset Subroutine” on page 14) subroutine, **chacl** (“chacl or fchacl Subroutine” on page 144) subroutine, **chmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine, “aclx\_convert Subroutine” on page 16.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## aclx\_gettypeinfo Subroutine

### Purpose

Retrieves the ACL characteristics given to an ACL type.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_gettypeinfo (Path, acl_type, buffer, buffer_sz)
char * Path;
acl_type_t  acl_type;
caddr_t  buffer;
size_t  * buffer_sz;
```

### Description

The **aclx\_gettypeinfo** subroutine helps obtain characteristics and capabilities of an ACL type on the file system. The buffer space provided by the caller is where the ACL type-related information is returned. If the length of this buffer is not enough to fit the characteristics for the ACL type requested, then **aclx\_gettypeinfo** returns an error and sets the *buffer\_len* field to the amount of buffer space needed.

### Parameters

<i>Path</i>	Specifies the path name of the file.
<i>acl_type</i>	ACL type for which the characteristics are sought.
<i>buffer</i>	Specifies the pointer to a buffer space, where the characteristics of <i>acl_type</i> for the file system is returned. The structure of data returned is ACL type-specific. Refer to the ACL type-specific documentation for more details.
<i>buffer_sz</i>	Points to an area that specifies the length of the buffer <i>buffer</i> in which the characteristics of <i>acl_type</i> are returned by the file system. This is an input/output parameter. If the length of the buffer provided is not sufficient to store all the ACL type characteristic information, then the file system returns an error and indicates the length of the buffer required in this variable. The length is specified in number of bytes.

### Return Values

On successful completion, the **aclx\_gettypeinfo** subroutine returns a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **aclx\_gettypeinfo** subroutine fails and the access control information for a file remains unchanged if one or more of the following is true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).

<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOSPC</b>	Buffer space provided is not enough to store all the <i>acl_type</i> characteristics of the file system.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

If Network File System (NFS) is installed on your system, the **acl\_gettypeinfo** subroutine can also fail if the following condition is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

Auditing Events: None

## Related Information

The “aclx\_get or aclx\_fget Subroutine” on page 17, “aclx\_put or aclx\_fput Subroutine” on page 25.

The **aclget** command, **aclput** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## aclx\_gettypes Subroutine

### Purpose

Retrieves the list of ACL types supported for the file system associated with the path provided.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_gettypes (Path, acl_type_list, acl_type_list_len)
char * Path;
acl_types_list_t * acl_type_list;
size_t * acl_type_list_len;
```

### Description

The **aclx\_gettypes** subroutine helps obtain the list of ACL types supported on the particular file system. A file system can implement policies to support one to many ACL types simultaneously. The first ACL type in the list is the default ACL type for the file system. This default ACL type is used in ACL conversions if the target ACL type is not supported on the file system. Each file system object in the file system is associated with only one piece of ACL data of a particular ACL type.

## Parameters

<i>Path</i>	Specifies the path name of the file system object within the file system for which the list of supported ACLs are being requested.
<i>acl_type_list</i>	Specifies the pointer to a buffer space, where the list of ACL types is returned. The size of this buffer is indicated using the <i>acl_type_list_len</i> argument in bytes.
<i>acl_type_list_len</i>	Pointer to a buffer that specifies the length of the buffer <i>acl_type_list</i> in which the list of ACLs is returned by the file system. This is an input/output parameter. If the length of the buffer is not sufficient to store all the ACL types, the file system returns an error and indicates the length of the buffer required in this same area. The length is specified in bytes. If the subroutine call is successful, this field contains the number of bytes of information stored in the <i>acl_type_list</i> buffer. This information can be used by the caller to get the number of ACL type entries returned.

## Return Values

On successful completion, the **aclx\_gettypes** subroutine returns a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **aclx\_gettypes** subroutine fails and the access control information for a file remains unchanged if one or more of the following is true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOSPC</b>	The <i>acl_type_list</i> buffer provided is not enough to store all the ACL types supported by this file system.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

If Network File System (NFS) is installed on your system, the **acl\_gettypes** subroutine can also fail if the following condition is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

Access Control: Caller must have search permission for all components of the *Path* prefix.

Auditing Events: None

## Related Information

The **aclget** command, **aclput** command.

## aclx\_print or aclx\_printStr Subroutine

### Purpose

Converts the binary access control information into nonbinary, readable format.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_print (acl_file, acl, acl_sz, acl_type, fs_obj_path, flags)
```

```
FILE * acl_file;
```

```
void * acl;
```

```
size_t  acl_sz;
```

```
acl_type_t  acl_type;
```

```
char * fs_obj_path;
```

```
int32_t  flags;
```

```
int aclx_printStr (str, str_sz, acl, acl_sz, acl_type, fs_obj_path, flags)
```

```
char * str;
```

```
size_t * str_sz;
```

```
void * acl;
```

```
size_t  acl_sz;
```

```
acl_type_t  acl_type;
```

```
char * fs_obj_path;
```

```
int32_t  flags;
```

### Description

The **aclx\_print** and **aclx\_printStr** subroutines print the access control information in a nonbinary, readable text format. These subroutines take the ACL information in binary format as input, convert it into text format, and print that text format output to either a file or a string. The **aclx\_print** subroutine prints the ACL text to the file specified by *acl\_file*. The **aclx\_printStr** subroutine prints the ACL text to *str*. The amount of space available in *str* is specified in *str\_sz*. If this memory is insufficient, the subroutine sets *str\_sz* to the needed amount of memory and returns an **ENOSPC** error.

### Parameters

<i>acl_file</i>	Points to the file into which the textual output is printed.
<i>str</i>	Points to the string into which the textual output should be printed.
<i>str_sz</i>	Indicates the amount of memory in bytes available in <i>str</i> . If the text representation of <i>acl</i> requires more space than <i>str_sz</i> , this subroutine updates the <i>str_sz</i> with the amount of memory required and fails by setting <b>errno</b> to <b>ENOSPC</b> .
<i>acl</i>	Points to a buffer which contains the binary ACL data that has to be printed. The size of this buffer is indicated by the <i>acl_sz</i> parameter.
<i>acl_sz</i>	Indicates the size of the buffer area passed through the <i>acl</i> parameter.
<i>acl_type</i>	Indicates the ACL type information of the <i>acl</i> . The ACL type is 64 bits in size and is unique on the system. If the given ACL type is not supported in the system, this function fails and <b>errno</b> is set to <b>EINVAL</b> .

<i>fs_obj_path</i>	File System Object Path for which the ACL data format and print are being requested. Gets information about the object (such as whether the object is a file or directory, who the owner is, and the associated group ID).
<i>flags</i>	Allows for control over the print operation. A value of <b>ACL_VERBOSE</b> indicates whether additional information has to be printed in text format in comments. This bit is set when the <b>aclget</b> command is issued with the <b>-v</b> (verbose) option.

## Return Values

On successful completion, the **aclx\_print** and **aclx\_printStr** subroutines return a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **aclx\_print** subroutine fails if one or more of the following is true:

**Note:** The errors in the following list occur only because **aclx\_print** calls the **fprintf** subroutine internally. For more information about these errors, refer to the **fprintf** subroutine.

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> flag is set for the file descriptor underlying the file specified by the <i>acl_file</i> parameter, and the process would be delayed in the write operation.
<b>EBADF</b>	The file descriptor underlying the file specified by the <i>acl_file</i> parameter is not a valid file descriptor open for writing.
<b>EFBIG</b>	An attempt was made to write to a file that exceeds the file size limit of this process or the maximum file size. For more information, refer to the <b>ulimit</b> subroutine.
<b>EINTR</b>	The write operation terminated because of a signal was received, and either no data was transferred or a partial transfer was not reported.
<b>EIO</b>	The process is a member of a background process group attempting to perform a write to its controlling terminal, the <b>TOSTOP</b> flag is set, the process is neither ignoring nor blocking the <b>SIGTTOU</b> signal, and the process group of the process has no parent process.
<b>ENOSPC</b>	No free space remains on the device that contains the file.
<b>ENOSPC</b>	Insufficient storage space is available.
<b>ENXIO</b>	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
<b>EPIPE</b>	An attempt was made to write to a pipe or first-in-first-out (FIFO) that is not open for reading by any process. A <b>SIGPIPE</b> signal is sent to the process.

The **aclx\_printStr** subroutine fails if the following is true:

<b>ENOSPC</b>	Input buffer size <i>strSz</i> is not sufficient to store the text representation of <i>acl</i> in <i>str</i> .
<b>ENOSPC</b>	Insufficient storage space is available. This error is returned by <b>sprintf</b> , which is called by the <b>aclx_printStr</b> subroutine internally.

The **aclx\_print** or **aclx\_printStr** subroutine fails if the following is true:

<b>EINVAL</b>	Invalid input parameter. The same error can be returned if an invalid <i>acl_type</i> is specified as input to this routine. This <b>errno</b> can also be returned if the <i>acl</i> is not of the type specified by <i>acl_type</i> .
---------------	---

## Related Information

The “printf, fprintf, sprintf, snprintf, wprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168, “aclx\_scan or aclx\_scanStr Subroutine” on page 27.

The **aclget** command, **aclput** command.



## aclx\_put or aclx\_fput Subroutine

### Purpose

Stores the access control information for a file system object.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_put (Path, ctl_flags, acl_type, acl, acl_sz, mode_info)
char * Path;
uint64_t  ctl_flags;
acl_type_t  acl_type;
void * acl;
size_t  acl_sz;
mode_t  mode_info;
```

```
int aclx_fput (FileDescriptor, ctl_flags, acl_type, acl, acl_sz, mode_info)
int FileDescriptor;
uint64_t  ctl_flags;
acl_type_t  acl_type;
void * acl;
size_t  acl_sz;
mode_t  mode_info;
```

### Description

The **aclx\_put** and **aclx\_fput** subroutines store the access control information for a file system object in the native ACL format. Native ACL format is the format as defined for the particular ACL type in the system. These subroutines are advanced versions of the **acl\_put** and **acl\_fput** subroutines and should be used instead of the older versions. The **aclx\_put** and **aclx\_fput** subroutines provide for more control for the user to interact with the underlying file system directly.

A caller specifies the ACL type in the *acl\_type* argument and passes the ACL information in the *acl* argument. The *acl\_sz* parameter indicates the size of the ACL data. The *ctl\_flags* parameter is a bitmask that allows for variation of **aclx\_put** requests.

The value provided to these subroutines can be obtained by invoking **aclx\_get** or **aclx\_fget** subroutines to copy or restore the access control information.

The **aclx\_put** and **aclx\_fput** subroutines can also be used to manage the special bits (such as SGID and SUID) in the mode word associated with the file system object. For example, you can set the **mode\_info** value to any special bit mask (as in the mode word defined for the file system), and a request can be made to set the same bits using the *ctl\_flags* argument. Note that special privileges (such as root) might be required to set these bits.

### Parameters

*Path* Specifies the path name of the file system object.

<i>FileDescriptor</i>	Specifies the file descriptor of an open file system object. This 64-bit sized bit mask provides control over the ACL retrieval. These bits are divided as follows: <b>Lower 16 bits</b> System-wide (nonphysical file-system-specific) ACL control flags <b>32 bits</b> Reserved. <b>Last 16 bits</b> Any physical file-system-defined options (that are specific to physical file system ACL implementation).
<i>ctl_flags</i>	Bit mask with the following system-wide flag values defined: <b>SET_MODE_S_BITS</b> Indicates that the <b>mode_info</b> value is set by the caller and the ACL put operation needs to consider this value while completing the ACL put operation. <b>SET_ACL</b> Indicates that the <i>acl</i> argument points to valid ACL data that needs to be considered while the ACL put operation is being performed.  <b>Note:</b> Both of the preceding values can be specified by the caller by ORing the two masks.
<i>acl_type</i>	Indicates the type of ACL to be associated with the file object. If the <i>acl_type</i> specified is not among the ACL types supported for the file system, then an error is returned.
<i>acl</i>	Points to a buffer where the ACL information exists. This ACL information is associated with the file system object specified. The size of this buffer is indicated by the <i>acl_sz</i> parameter.
<i>acl_sz</i>	Indicates the size of the ACL information sent through the <i>acl</i> parameter.
<i>mode_info</i>	This value indicates any mode word information that needs to be set for the file system object in question as part of this ACL put operation. When mode bits are being altered by specifying the <b>SET_MODE_S_BITS</b> flag (in <i>ctl_flags</i> ) ACL put operation fails if the caller does not have the required privileges.

## Return Values

On successful completion, the **aclx\_put** and **aclx\_fput** subroutines return a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **aclx\_put** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.

The **aclx\_fput** subroutine fails and the file permissions remain unchanged if the following is true:

**EBADF**                    The *FileDescriptor* parameter is not a valid file descriptor.

The **aclx\_put** or **aclx\_fput** subroutine fails if one or more of the following is true:

**EINVAL**                    Invalid input parameter. The same error can be returned if an invalid *acl\_type* is specified as input to this routine.  
**EIO**                        An I/O error occurred during the operation.  
**EROFS**                    The named file resides on a read-only file system.

If Network File System (NFS) is installed on your system, the **acl\_put** and **acl\_fput** subroutines can also fail if the following condition is true:

**ETIMEDOUT**                The connection timed out.

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

Auditing Events:

Event	Information
<b>chacl</b>	<i>Path</i> -based event
<b>fchacl</b>	<i>FileDescriptor</i> -based event

## Related Information

The **acl\_chg** (“*acl\_chg* or *acl\_fchg* Subroutine” on page 8) subroutine, **acl\_get** (“*acl\_get* or *acl\_fget* Subroutine” on page 10) subroutine, **acl\_set** (“*acl\_set* or *acl\_fset* Subroutine” on page 14) subroutine, **chacl** (“*chacl* or *fchacl* Subroutine” on page 144) subroutine, **chmod** (“*chmod* or *fchmod* Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## aclx\_scan or aclx\_scanStr Subroutine

### Purpose

Reads the access control information that is in nonbinary, readable text format, and converts it into ACL type-specific native format binary ACL data.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
```

```
int aclx_scan (acl_file, acl, acl_sz, acl_type, err_file)  
FILE * acl_file;  
void * acl;
```

```

size_t * acl_sz;
acl_type_t acl_type;
FILE * err_file;

```

```

int aclx_scanStr (str, acl, acl_sz, acl_type)
char * str;
void * acl;
size_t * acl_sz;
acl_type_t acl_type;

```

## Description

The **aclx\_scan** and **aclx\_scanStr** subroutines read the access control information from the input given in nonbinary, readable text format and return a binary ACL data in the ACL type-specific native format. The **aclx\_scan** subroutine provides the ACL data text in the file specified by *acl\_file*. In the case of **aclx\_scanStr**, the ACL data text is provided in the string pointed to by *str*. When the *err\_file* parameter is not Null, it points to a file to which any error messages are written out by the **aclx\_scan** subroutine in case of syntax errors in the input ACL data. The errors can occur if the syntax of the input text data does not adhere to the required ACL type-specific data specifications.

## Parameters

<i>acl_file</i>	Points to the file from which the ACL text output is read.
<i>str</i>	Points to the string from which the ACL text output is printed.
<i>acl</i>	Points to a buffer in which the binary ACL data has to be stored. The amount of memory available in this buffer is indicated by the <i>acl_sz</i> parameter.
<i>acl_sz</i>	Indicates the amount of memory, in bytes, available in the <i>acl</i> parameter.
<i>acl_type</i>	Indicates the ACL type information of the <i>acl</i> . The ACL type is 64 bits in size and is unique on the system. If the given ACL type is not supported in the system, this function fails and <b>errno</b> is set to <b>EINVAL</b> .
<i>err_file</i>	File pointer to an error file. When this pointer is supplied, the subroutines write out any errors in the syntax/composition of the ACL input data.

## Return Values

On successful completion, the **aclx\_scan** and **aclx\_scanStr** subroutines return a value of 0. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **aclx\_scan** subroutine fails if one or more of the following is true:

**Note:** The errors in the following list occur only because **aclx\_scan** calls the **fscanf** subroutine internally. For more information about these errors, refer to the **fscanf** subroutine.

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> flag is set for the file descriptor underlying the file specified by the <i>acl_file</i> parameter, and the process would be delayed in the write operation.
<b>EBADF</b>	The file descriptor underlying the file specified by the <i>acl_file</i> parameter is not a valid file descriptor open for writing.
<b>EINTR</b>	The write operation terminated because of a signal was received, and either no data was transferred or a partial transfer was not reported.
<b>EIO</b>	The process is a member of a background process group attempting to perform a write to its controlling terminal, the <b>TOSTOP</b> flag is set, the process is neither ignoring nor blocking the <b>SIGTTOU</b> signal, and the process group of the process has no parent process.
<b>ENOSPC</b>	Insufficient storage space is available.

The **aclx\_scanStr** subroutine fails if the following is true:

**ENOSPC**                    Insufficient storage space is available. This error is returned by **sscanf**, which is called by the **aclx\_scanStr** subroutine internally.

The **aclx\_scan** or **aclx\_scanStr** subroutine fails if the following is true:

**EINVAL**                    Invalid input parameter. The same error can be returned if an invalid *acl\_type* is specified as input to this routine. This **errno** can also be returned if the text ACL given in the input/file string is not of the type specified by *acl\_type*.

## Related Information

The “aclx\_print or aclx\_printStr Subroutine” on page 23, fscanf Subroutine.

The **aclget** command, **aclput** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## acos, acosf, or acosl Subroutine

### Purpose

Computes the inverse cosine of a given value.

### Syntax

```
#include <math.h>
```

```
float acosf (x)  
float x;
```

```
long double acosl (x)  
long double x;
```

```
double acos (x)  
double x;
```

### Description

The **acosf**, **acosl**, and **acos** subroutines compute the principal value of the arc cosine of the *x* parameter. The value of *x* should be in the range [-1,1].

An application wishing to check for error situations should set the **errno** global variable to zero and call **fetestexcept(FE\_ALL\_EXCEPT)** before calling these functions. On return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x*                            Specifies the value to be computed.

## Return Values

Upon successful completion, these subroutines return the arc cosine of  $x$ , in the range  $[0, \pi]$  radians.

For finite values of  $x$  not in the range  $[-1, 1]$ , a domain error occurs, and a NaN is returned.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $+1$ , 0 is returned.

If  $x$  is  $\pm\text{Inf}$ , a domain error occurs, and a NaN is returned.

## Related Information

The “acosh, acoshf, or acoshl Subroutine.”

`math.h` in *AIX 5L Version 5.3 Files Reference*.

---

## acosh, acoshf, or acoshl Subroutine

### Purpose

Computes the inverse hyperbolic cosine.

### Syntax

```
#include <math.h>
```

```
float acoshf (x)
float x;
```

```
long double acoshl (x)
long double x;
```

```
double acosh (x)
double x;
```

### Description

The **acoshf**, or **acoshl** subroutine computes the inverse hyperbolic cosine of the  $x$  parameter.

The **acosh** subroutine returns the hyperbolic arc cosine specified by the  $x$  parameter, in the range 1 to the **+HUGE\_VAL** value.

An application wishing to check for error situations should set **errno** to zero and call **fetestexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if the **errno** global variable is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

## Return Values

Upon successful completion, the **acoshf**, or **acoshl** subroutine returns the inverse hyperbolic cosine of the given argument.

For finite values of  $x < 1$ , a domain error occurs, and a NaN is returned.

If *x* is NaN, a NaN is returned.

If *x* is +1, 0 is returned.

If *x* is +Inf, +Inf is returned.

If *x* is -Inf, a domain error occurs, and a NaN is returned.

## Error Codes

The **acosh** subroutine returns **NaNQ** (not-a-number) and sets **errno** to **EDOM** if the *x* parameter is less than the value of 1.

## Related Information

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## addproj Subroutine

### Purpose

Adds an API-based project definition to the kernel project registry.

### Library

The **libaacct.a** library.

### Syntax

```
<sys/aacct.h>
```

```
addproj(struct project *)
```

### Description

The **addproj** subroutine defines the application-based project definition to the kernel repository. An application can assign a project defined in this way using the **proj\_execve** system call.

Projects that are added this way are marked as being specified by applications so that they do not overlap with system administrator-specified projects defined using the **projctl** command. The **PROJFLAG\_API** flag is turned on in the structure project to indicate that the project definition was added by an application.

Projects added by a system administrator using the **projctl** command are flagged as being derived from the local or LDAP-based project repositories by the **PROJFLAGS\_LDAP** or **PROJFLAGS\_PDF** flag. If one of these flags is specified, the **addproj** subroutine fails with **EPERM**.

The **getproj** routine can be used to determine the origin of a loaded project.

The **addproj** validates the input project number to ensure that it is within the expected range of 0x00000001 - 0x00ffffff. It also validates that the project name is a POSIX compliant alphanumeric character string. If any invalid input is found **errno** will be set to **EINVAL** and the **addproj** subroutine returns -1.

### Parameters

*project*                      Points to a project structure that holds the definition of the project to be added.

## Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid Project Name / Number or the passed pointer is NULL
<b>EEXIST</b>	Project Definition exists
<b>EPERM</b>	Permission Denied, not a privileged user

## Related Information

The “addprojdb Subroutine,” “chprojattr Subroutine” on page 158, “getproj Subroutine” on page 418, “getprojs Subroutine” on page 420, rmproj Subroutine.

---

## addprojdb Subroutine

### Purpose

Adds a project definition to the specified project database.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
addprojdb(void *handle, struct project *project, char *comment)
```

### Description

The **addprojdb** subroutine appends the project definition stored in the struct *project* variable into the project database named by the *handle* parameter. The project database must be initialized before calling this subroutine. The **projdballoc** subroutine is provided for this purpose. This routine verifies whether the supplied project definition already exists. If it does exist, the **addprojdb** subroutine sets errno to **EEXIST** and returns -1.

The **addprojdb** subroutine validates the input project number to ensure that it is within the expected range 0x00000001 - 0x00ffffff and validates that the project name is a POSIX-compliant alphanumeric character string. If any invalid input is found, the **addprojdb** subroutine sets errno to **EINVAL** and returns -1.

If the user does not have privilege to add an entry to project database, the **addprojdb** subroutine sets errno to **EACCES** and returns -1.

There is an internal state (that is, the current project) associated with the project database. When the project database is initialized, the current project is the first project in the database. The **addprojdb** subroutine appends the specified project to the end of the database. It advances the current project assignment to the next project in the database, which is the end of the project data base. At this point, a call to the **getnextprojdb** subroutine would fail, because there are no additional project definitions. To read



the project definition that was just added, use the **getprojdb** subroutine. To read other projects, first call **getfirstprojdb** subroutine to reset the internal current project assignment so that subsequent reads can be performed.

The format of the records added to the project database are given as follows:

```
ProjectName:ProjectNumber:AggregationStatus:Comment::
```

Example:

```
Biology:4756:no:Project Created by project1 command::
```

## Parameters

<i>handle</i>	Pointer to project database handle
<i>project</i>	Pointer to a project structure that holds the definition of the project to be added
<i>comment</i>	Pointer to a character string that holds the comments about the project

## Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid project name or number, or the passed pointer is NULL.
<b>EEXIST</b>	Project definition already exists.
<b>EPERM</b>	Permission denied. The user is not a privileged user.

## Related Information

The “addproj Subroutine” on page 31, “chprojattrdb Subroutine” on page 159, “getfirstprojdb Subroutine” on page 368, “getnextprojdb Subroutine” on page 396, “getprojdb Subroutine” on page 419, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, “projdbfree Subroutine” on page 1181, rmprojdb Subroutine.

---

## addssys Subroutine

### Purpose

Adds the **SRCsubsys** record to the subsystem object class.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <sys/srcobj.h>
#include <spc.h>
```

```
int addssys ( SRCSubsystem )
struct SRCsubsys *SRCSubsystem;
```

## Description

The **addssys** subroutine adds a record to the subsystem object class. You must call the **defssys** subroutine to initialize the *SRCSubsystem* buffer before your application program uses the **SRCsubsys** structure. The **SRCsubsys** structure is defined in the */usr/include/sys/srcobj.h* file.

The executable running with this subroutine must be running with the group system.

## Parameters

*SRCSubsystem*                      A pointer to the **SRCsubsys** structure.

## Return Values

Upon successful completion, the **addssys** subroutine returns a value of 0. Otherwise, it returns a value of -1 and the **odmerrno** variable is set to indicate the error, or an SRC error code is returned.

## Error Codes

The **addssys** subroutine fails if one or more of the following are true:

<b>SRC_BADFSIG</b>	Invalid stop force signal.
<b>SRC_BADNSIG</b>	Invalid stop normal signal.
<b>SRC_CMDARG2BIG</b>	Command arguments too long.
<b>SRC_GRPNAM2BIG</b>	Group name too long.
<b>SRC_NOCONTACT</b>	Contact not signal, sockets, or message queue.
<b>SRC_NONAME</b>	No subsystem name specified.
<b>SRC_NOPATH</b>	No subsystem path specified.
<b>SRC_PATH2BIG</b>	Subsystem path too long.
<b>SRC_STDERR2BIG</b>	stderr path too long.
<b>SRC_STDIN2BIG</b>	stdin path too long.
<b>SRC_STDOUT2BIG</b>	stdout path too long.
<b>SRC_SUBEXIST</b>	New subsystem name already on file.
<b>SRC_SUBSYS2BIG</b>	Subsystem name too long.
<b>SRC_SYNEXIST</b>	New subsystem synonym name already on file.
<b>SRC_SYN2BIG</b>	Synonym name too long.

## Security

Privilege Control: This command has the Trusted Path attribute. It has the following kernel privilege:

SET\_PROC\_AUDIT  
Files Accessed:

Mode	File
644	<i>/etc/objrepos/SRCsubsys</i>

Auditing Events:

If the auditing subsystem has been properly configured and is enabled, the **addssys** subroutine generates the following audit record (event) each time the subroutine is executed:

Event	Information
<b>SRC_addssys</b>	Lists the SRCsubsys records added.

See "Setting Up Auditing" in *Security* for details about selecting and grouping audit events, and configuring audit event data collection.

## Files

<b>/etc/objrepos/SRCsubsys</b>	SRC Subsystem Configuration object class.
<b>/dev/SRC</b>	Specifies the <b>AF_UNIX</b> socket file.
<b>/dev/SRC-unix</b>	Specifies the location for temporary socket files.
<b>/usr/include/spc.h</b>	Defines external interfaces provided by the SRC subroutines.
<b>/usr/include/sys/srcobj.h</b>	Defines object structures used by the SRC.

## Related Information

The **chssys** ("chssys Subroutine" on page 162) subroutine, **defssys** ("defssys Subroutine" on page 211) subroutine, **delssys** ("delssys Subroutine" on page 212) subroutine.

The **auditpr** command, **chssys** command, **mkssys** command, **rmssys** command.

Auditing Overview ("audit Subroutine" on page 98) and System Resource Controller in *Operating system and device management*.

Defining Your Subsystem to the SRC, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

List of SRC Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## adjtime Subroutine

### Purpose

Corrects the time to allow synchronization of the system clock.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/time.h>
int adjtime ( Delta, Olddelta)
struct timeval *Delta;
struct timeval *Olddelta;
```

### Description

The **adjtime** subroutine makes small adjustments to the system time, as returned by the **gettimeofday** subroutine, advancing or retarding it by the time specified by the *Delta* parameter of the **timeval** structure. If the *Delta* parameter is negative, the clock is slowed down by periodically subtracting a small amount from it until the correction is complete. If the *Delta* parameter is positive, a small amount is periodically added to the clock until the correction is complete. The skew used to perform the correction is generally ten percent. If the clock is sampled frequently enough, an application program can see time apparently jump backwards. For information on a way to avoid this, see "gettimeofday, settimeofday, or ftime

Subroutine” on page 445. A time correction from an earlier call to the **adjtime** subroutine may not be finished when the **adjtime** subroutine is called again. If the *Olddelta* parameter is nonzero, then the structure pointed to will contain, upon return, the number of microseconds still to be corrected from the earlier call.

This call may be used by time servers that synchronize the clocks of computers in a local area network. Such time servers would slow down the clocks of some machines and speed up the clocks of others to bring them to the average network time.

The **adjtime** subroutine is restricted to the users with root user authority.

## Parameters

*Delta* Specifies the amount of time to be altered.  
*Olddelta* Contains the number of microseconds still to be corrected from an earlier call.

## Return Values

A return value of 0 indicates that the **adjtime** subroutine succeeded. A return value of -1 indicates that an error occurred, and **errno** is set to indicate the error.

## Error Codes

The **adjtime** subroutine fails if the following are true:

<b>EFAULT</b>	An argument address referenced invalid memory.
<b>EPERM</b>	The process's effective user ID does not have root user authority.

---

## agg\_proc\_stat, agg\_lpar\_stat, agg\_arm\_stat, or free\_agg\_list Subroutine

### Purpose

Aggregate advanced accounting data.

### Library

The **libaacct.a** library.

### Syntax

```
#define <sys/aacct.h>
int agg_arm_stat(tran_list, arm_list);
struct aacct_tran_rec *tran_list
struct agg_arm_stat **arm_list
int agg_proc_stat(sortcrit1, sortcrit2, sortcrit3, sortcrit4, tran_list, proc_list);
int sortcrit1, sortcrit2, sortcrit3, sortcrit4
struct aacct_tran_rec *tran_list
struct agg_proc_stat **proc_list
int agg_lpar_stat(l_type, *tran_list, l_list);
int l_type
struct aacct_tran_rec *tran_list
union agg_lpar_rec *l_list
void free_agg_list(list);
void *list
```

## Description

The **agg\_proc\_stat**, **agg\_lpar\_stat**, and **agg\_arm\_stat** subroutines return a linked list of aggregated transaction records for process, LPAR, and ARM, respectively.

The **agg\_proc\_stat** subroutine performs the process record aggregation based on the criterion values passed as input parameters. The aggregated process transaction records are sorted based on the sorting criteria values *sortcrit1*, *sortcrit2*, *sortcrit3*, and *sortcrit4*. These four can be one of the following values defined in the **sys/aacct.h** file:

- **CRIT\_UID**
- **CRIT\_GID**
- **CRIT\_PROJ**
- **CRIT\_CMD**
- **CRIT\_NONE**

The order of their usage determines the sorting order applied to the retrieved aggregated list of process transaction records. For example, the sort criteria values of **PROJ\_GID**, **PROJ\_PROJ**, **PROJ\_UID**, **PROJ\_NONE** first sorts the aggregated list on group IDs, which are further sorted based on project IDs, followed by another level of sorting based on user IDs.

Some of the process transaction records (of type **TRID\_agg\_proc**) cannot be aggregated based on group IDs and command names. For such records, **agg\_proc\_stat** returns an asterisk (\*) character as the command name and a value of -2 as the group ID. This indicates to the caller that these records cannot be aggregated.

If the aggregation is not necessary on a specific criteria, **agg\_proc\_stat** returns a value of -1 in the respective field. For example, if the aggregation is not necessary on the group ID (**CRIT\_GID**), the retrieved list of aggregation records has a value of -1 filled in the group ID fields.

The **agg\_lpar\_stat** retrieves an aggregated list of LPAR transaction records. Because there are several types of LPAR transaction records, the caller must specify the type of LPAR transaction record that is to be aggregated. The transaction record type can be one of the following values, defined in the **sys/aacct.h** file:

- **AGG\_CPUMEM**
- **AGG\_FILESYS**
- **AGG\_NETIF**
- **AGG\_DISK**
- **AGG\_VTARGET**
- **AGG\_VCLIENT**

The **agg\_lpar\_stat** subroutine uses a union argument of type **struct agg\_lpar\_rec**. For this argument, the caller must provide the address of the linked list to which the aggregated records should be returned.

The **agg\_arm\_list** retrieves an aggregated list of ARM transaction records from the list of transaction records provided as input. The aggregated transaction records are returned to the caller through the structure pointer of type **struct agg\_arm\_stat**.

The **free\_agg\_list** subroutine frees the memory allocated to the aggregated records returned by the **agg\_proc\_stat**, **agg\_lpar\_stat**, or **agg\_arm\_stat** subroutine.

## Parameters

*arm\_list*

Pointer to the linked list of **struct agg\_arm\_stat** nodes to be returned.

<i>_list</i>	Pointer to <b>union agg_lpar_rec</b> address to which the aggregated LPAR records are returned.
<i>_type</i>	Integer value that represents the type of LPAR resource to be aggregated.
<i>list</i>	Pointer to the aggregated list to be freed.
<i>proc_list</i>	Pointer to the linked list of <b>struct agg_proc_stat</b> nodes to be returned.
<i>sortcrit1, sortcrit2, sortcrit3, sortcrit4</i>	Integer values that represent the sorting criteria to be passed to <b>agg_proc_stat</b> .
<i>tran_list</i>	Pointer to the input list of transaction records

## Security

No restrictions. Any user can call this function.

## Return Values

0	The call to the subroutine was successful.
-1	The call to the subroutine failed.

## Error Codes

<b>EINVAL</b>	The passed pointer is NULL.
<b>ENOMEM</b>	Insufficient memory.
<b>EPERM</b>	Permission denied. Unable to read the data file.

## Related Information

The “buildproclist Subroutine” on page 125, “buildtranlist or freetranlist Subroutine” on page 126, “getproclist, getlparlist, or getarmlist Subroutine” on page 414.

Understanding the Advanced Accounting Subsystem.

---

## aio\_cancel or aio\_cancel64 Subroutine

The **aio\_cancel** or **aio\_cancel64** subroutine includes information for the POSIX AIO **aio\_cancel** subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO **aio\_cancel** subroutine.

### POSIX AIO aio\_cancel Subroutine

## Purpose

Cancels one or more outstanding asynchronous I/O requests.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>
```

```
int aio_cancel (fildes, aiocbp)
int fildes;
struct aiocb *aiocbp;
```

## Description

The **aiocancel** subroutine cancels one or more asynchronous I/O requests currently outstanding against the *filides* parameter. The *aiocbp* parameter points to the asynchronous I/O control block for a particular request to be canceled. If *aiocbp* is NULL, all outstanding cancelable asynchronous I/O requests against *filides* are canceled.

Normal asynchronous notification occurs for asynchronous I/O operations that are successfully canceled. If there are requests that cannot be canceled, the normal asynchronous completion process takes place for those requests when they are completed.

For requested operations that are successfully canceled, the associated error status is set to **ECANCELED**, and a -1 is returned. For requested operations that are not successfully canceled, the *aiocbp* parameter is not modified by the **aiocancel** subroutine.

If *aiocbp* is not NULL, and if *filides* does not have the same value as the file descriptor with which the asynchronous operation was initiated, unspecified results occur.

The implementation of the subroutine defines which operations are cancelable.

## Parameters

*filides*                Identifies the object to which the outstanding asynchronous I/O requests were originally queued.  
*aiocbp*                Points to the **aiocb** structure associated with the I/O operation.

## aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
int                    aio_filides
off_t                  aio_offset
char                   *aio_buf
size_t                 aio_nbytes
int                    aio_reqprio
struct sigevent        aio_sigevent
int                    aio_lio_opcode
```

## Execution Environment

The **aiocancel** and **aiocancel64** subroutines can be called from the process environment only.

## Return Values

The **aiocancel** subroutine returns **AIO\_CANCELED** to the calling process if the requested operation(s) were canceled. **AIO\_NOTCANCELED** is returned if at least one of the requested operations cannot be canceled because it is in progress. In this case, the state of the other operations, referenced in the call to **aiocancel** is not indicated by the return value of **aiocancel**. The application may determine the state of affairs for these operations by using the **aio\_error** subroutine. **AIO\_ALLDONE** is returned if all of the operations are completed. Otherwise, the subroutine returns -1 and sets the **errno** global variable to indicate the error.

## Error Codes

**EBADF**                The *filides* parameter is not a valid file descriptor.

## Related Information

“**aio\_error** or **aio\_error64** Subroutine” on page 42, “**aio\_nwait** Subroutine” on page 47, “**aio\_nwait\_timeout** Subroutine” on page 49, “**aio\_read** or **aio\_read64** Subroutine” on page 50, “**aio\_return** or **aio\_return64**”

Subroutine” on page 55, “aio\_suspend or aio\_suspend64 Subroutine” on page 58, “aio\_write or aio\_write64 Subroutine” on page 61, and “lio\_listio or lio\_listio64 Subroutine” on page 720.

The Asynchronous I/O Subsystem and Communications I/O Subsystem in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

## Legacy AIO `aio_cancel` Subroutine

### Purpose

Cancels one or more outstanding asynchronous I/O requests.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <aio.h>
```

```
aio_cancel ( FileDescriptor, aiocbp)
int FileDescriptor;
struct aiocb *aiocbp;
```

```
aio_cancel64 ( FileDescriptor, aiocbp)
int FileDescriptor;
struct aiocb64 *aiocbp;
```

### Description

The **aio\_cancel** subroutine attempts to cancel one or more outstanding asynchronous I/O requests issued on the file associated with the *FileDescriptor* parameter. If the pointer to the **aio control block (aiocb)** structure (the *aiocbp* parameter) is not null, then an attempt is made to cancel the I/O request associated with this **aiocb**. The *aiocbp* parameter used by the thread calling **aio\_cancel** must have had its request initiated by this same thread. Otherwise, a -1 is returned and **errno** is set to EINVAL. However, if the *aiocbp* parameter is null, then an attempt is made to cancel all outstanding asynchronous I/O requests associated with the *FileDescriptor* parameter without regard to the initiating thread.

The **aio\_cancel64** subroutine is similar to the **aio\_cancel** subroutine except that it attempts to cancel outstanding large file enabled asynchronous I/O requests. Large file enabled asynchronous I/O requests make use of the **aiocb64** structure instead of the **aiocb** structure. The **aiocb64** structure allows asynchronous I/O requests to specify offsets in excess of OFF\_MAX (2 gigbytes minus 1).

In the large file enabled programming environment, **aio\_cancel** is redefined to be **aio\_cancel64**.

When an I/O request is canceled, the **aio\_error** (“aio\_error or aio\_error64 Subroutine” on page 42) subroutine called with the handle to the corresponding **aiocb** structure returns **ECANCELED**.

**Note:** The `_AIO_AIX_SOURCE` macro used in **aio.h** must be defined when using **aio.h** to compile an aio application with the Legacy AIO function definitions. The default compile using the **aio.h** file is for an application with the POSIX AIO definitions. In the source file enter:

```
#define _AIO_AIX_SOURCE
#include <sys/aio.h>
```



or, on the command line when compiling enter:  
->xlc ... -D\_AIO\_AIX\_SOURCE ... legacy\_aio\_program.c

## Parameters

*FileDescriptor* Identifies the object to which the outstanding asynchronous I/O requests were originally queued.  
*aiocbp* Points to the **aiocb** structure associated with the I/O operation.

## aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
struct aiocb
{
    int          aio_whence;
    off_t        aio_offset;
    char         *aio_buf;
    ssize_t      aio_return;
    int          aio_errno;
    size_t       aio_nbytes;
    union {
        int      reqprio;
        struct {
            int  version:8;
            int  priority:8;
            int  cache_hint:16;
        } ext;
    } aio_ul;
    int          aio_flag;
    int          aio_iocpfd;
    aio_handle_t aio_handle;
}

#define aio_reqprio      aio_ul.reqprio
#define aio_version     aio_ul.ext.version
#define aio_priority    aio_ul.ext.priority
#define aio_cache_hint  aio_ul.ext.cache_hint
```

## Execution Environment

The **aio\_cancel** and **aio\_cancel64** subroutines can be called from the process environment only.

## Return Values

**AIO\_CANCELED** Indicates that all of the asynchronous I/O requests were canceled successfully. The **aio\_error** subroutine call with the handle to the **aiocb** structure of the request will return **ECANCELED**.

**AIO\_NOTCANCELED** Indicates that the **aio\_cancel** subroutine did not cancel one or more outstanding I/O requests. This may happen if an I/O request is already in progress. The corresponding error status of the I/O request is not modified.

**AIO\_ALLDONE** Indicates that none of the I/O requests is in the queue or in progress.

**-1** Indicates that the subroutine was not successful. Sets the **errno** global variable to identify the error.

A return code can be set to the following **errno** value:

**EBADF** Indicates that the *FileDescriptor* parameter is not valid.

## Related Information

“aio\_error or aio\_error64 Subroutine,” “aio\_nwait Subroutine” on page 47, “aio\_nwait\_timeout Subroutine” on page 49, “aio\_read or aio\_read64 Subroutine” on page 50, “aio\_return or aio\_return64 Subroutine” on page 55, “aio\_suspend or aio\_suspend64 Subroutine” on page 58, and “aio\_write or aio\_write64 Subroutine” on page 61, “lio\_listio or lio\_listio64 Subroutine” on page 720.

The Asynchronous I/O Subsystem and Communications I/O Subsystem in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## aio\_error or aio\_error64 Subroutine

The **aio\_error** or **aio\_error64** subroutine includes information for the POSIX AIO **aio\_error** subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO **aio\_error** subroutine.

### POSIX AIO aio\_error Subroutine

## Purpose

Retrieves error status for an asynchronous I/O operation.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>

int aio_error (aiocbp)
const struct aiocb *aiocbp;
```

## Description

The **aio\_error** subroutine returns the error status associated with the **aiocb** structure. This structure is referenced by the *aiocbp* parameter. The error status for an asynchronous I/O operation is the synchronous I/O **errno** value that would be set by the corresponding **read**, **write**, or **fsync** subroutine. If the subroutine has not yet completed, the error status is equal to **EINPROGRESS**.

## Parameters

*aiocbp*                      Points to the **aiocb** structure associated with the I/O operation.

## aiocb Structure

The **aiocb** structure is defined in the **/usr/include/aio.h** file and contains the following members:

```
int                      aio_fildes
off_t                    aio_offset
char                    *aio_buf
size_t                  aio_nbytes
int                      aio_reqprio
struct sigevent        aio_sigevent
int                      aio_lio_opcode
```

## Execution Environment

The `aio_error` and `aio_error64` subroutines can be called from the process environment only.

## Return Values

If the asynchronous I/O operation has completed successfully, the `aio_error` subroutine returns a 0. If unsuccessful, the error status (as described for the `read`, `write`, and `fsync` subroutines) is returned. If the asynchronous I/O operation has not yet completed, `EINPROGRESS` is returned.

## Error Codes

**EINVAL**            The `aiocbp` parameter does not refer to an asynchronous operation whose return status has not yet been retrieved.

## Related Information

“`aio_cancel` or `aio_cancel64` Subroutine” on page 38, “`aio_fsync` Subroutine” on page 45, “`aio_nwait` Subroutine” on page 47, “`aio_nwait_timeout` Subroutine” on page 49, “`aio_read` or `aio_read64` Subroutine” on page 50, “`aio_return` or `aio_return64` Subroutine” on page 55, “`aio_write` or `aio_write64` Subroutine” on page 61, “`close` Subroutine” on page 175, “`exec: execl, execl, execlp, execv, execve, execvp, or exect` Subroutine” on page 236, “`exit, atexit, unatexit, _exit, or _Exit` Subroutine” on page 243, “`fork, f_fork, or vfork` Subroutine” on page 291, “`fsync` or `fsync_range` Subroutine” on page 321, “`lio_listio` or `lio_listio64` Subroutine” on page 720, and “`lseek, llseek` or `lseek64` Subroutine” on page 765.

`read`, `readx`, `readv`, `readvx`, or `pread` Subroutine and `write`, `writex`, `writev`, `writevx` or `pwrite` Subroutines in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

## Legacy AIO `aio_error` Subroutine

## Purpose

Retrieves the error status of an asynchronous I/O request.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>
```

```
int  
aio_error(handle)  
aio_handle_t handle;
```

```
int aio_error64(handle)  
aio_handle_t handle;
```

## Description

The `aio_error` subroutine retrieves the error status of the asynchronous request associated with the `handle` parameter. The error status is the `errno` value that would be set by the corresponding I/O operation. The error status is `EINPROG` if the I/O operation is still in progress.

The `aio_error64` subroutine is similar to the `aio_error` subroutine except that it retrieves the error status associated with an `aiocb64` control block.

**Note:** The `_AIO_AIX_SOURCE` macro used in `aio.h` must be defined when using `aio.h` to compile an aio application with the Legacy AIO function definitions. The default compile using the `aio.h` file is for an application with the POSIX AIO definitions. In the source file enter:

```
#define _AIO_AIX_SOURCE
#include <sys/aio.h>
```

or, on the command line when compiling enter:

```
->xlc ... -D_AIO_AIX_SOURCE ... legacy_aio_program.c
```

## Parameters

*handle* The handle field of an **aio control block** (`aioCb` or `aioCb64`) structure set by a previous call of the `aio_read`, `aio_read64`, `aio_write`, `aio_write64`, `lio_listio`, `aio_listio64` subroutine. If a random memory location is passed in, random results are returned.

## aioCb Structure

The `aioCb` structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
struct aioCb
{
    int             aio_whence;
    off_t           aio_offset;
    char            *aio_buf;
    ssize_t         aio_return;
    int             aio_errno;
    size_t          aio_nbytes;
    union {
        int         reqprio;
        struct {
            int     version:8;
            int     priority:8;
            int     cache_hint:16;
        } ext;
    } aio_u1;
    int             aio_flag;
    int             aio_iocpfd;
    aio_handle_t    aio_handle;
}

#define aio_reqprio      aio_u1.reqprio
#define aio_version     aio_u1.ext.version
#define aio_priority    aio_u1.ext.priority
#define aio_cache_hint  aio_u1.ext.cache_hint
```

## Execution Environment

The `aio_error` and `aio_error64` subroutines can be called from the process environment only.

## Return Values

**0** Indicates that the operation completed successfully.

**ECANCELED** Indicates that the I/O request was canceled due to an `aio_cancel` subroutine call.

**EINPROG** Indicates that the I/O request has not completed.

An **errno** value described in the **aio\_read** (“aio\_read or aio\_read64 Subroutine” on page 50), **aio\_write** (“aio\_write or aio\_write64 Subroutine” on page 61), and **lio\_listio** (“lio\_listio or lio\_listio64 Subroutine” on page 720) subroutines: Indicates that the operation was not queued successfully. For example, if the **aio\_read** subroutine is called with an unusable file descriptor, it (**aio\_read**) returns a value of -1 and sets the **errno** global variable to **EBADF**. A subsequent call of the **aio\_error** subroutine with the handle of the unsuccessful **aio control block (aiocb)** structure returns **EBADF**.

An **errno** value of the corresponding I/O operation: Indicates that the operation was initiated successfully, but the actual I/O operation was unsuccessful. For example, calling the **aio\_write** subroutine on a file located in a full file system returns a value of 0, which indicates the request was queued successfully. However, when the I/O operation is complete (that is, when the **aio\_error** subroutine no longer returns **EINPROG**), the **aio\_error** subroutine returns **ENOSPC**. This indicates that the I/O was unsuccessful.

## Related Information

“aio\_cancel or aio\_cancel64 Subroutine” on page 38, “aio\_read or aio\_read64 Subroutine” on page 50, “aio\_nwait Subroutine” on page 47, “aio\_nwait\_timeout Subroutine” on page 49, “aio\_return or aio\_return64 Subroutine” on page 55, “aio\_suspend or aio\_suspend64 Subroutine” on page 58, “aio\_write or aio\_write64 Subroutine” on page 61, “lio\_listio or lio\_listio64 Subroutine” on page 720, and “lio\_listio or lio\_listio64 Subroutine” on page 720.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## aio\_fsync Subroutine

### Purpose

Synchronizes asynchronous files.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <aio.h>

int aio_fsync (op, aiocbp)
int op;
struct aiocb *aiocbp;
```

### Description

The **aio\_fsync** subroutine asynchronously forces all I/O operations to the synchronized I/O completion state. The function call returns when the synchronization request has been initiated or queued to the file or device (even when the data cannot be synchronized immediately).

If the *op* parameter is set to **O\_DSYNC**, all currently queued I/O operations are completed as if by a call to the **fdatasync** subroutine. If the *op* parameter is set to **O\_SYNC**, all currently queued I/O operations are

completed as if by a call to the **fsync** subroutine. If the **aio\_fsync** subroutine fails, or if the operation queued by **aio\_fsync** fails, outstanding I/O operations are not guaranteed to be completed.

If **aio\_fsync** succeeds, it is only the I/O that was queued at the time of the call to **aio\_fsync** that is guaranteed to be forced to the relevant completion state. The completion of subsequent I/O on the file descriptor is not guaranteed to be completed in a synchronized fashion.

The *aiocbp* parameter refers to an asynchronous I/O control block. The *aiocbp* value can be used as an argument to the **aio\_error** and **aio\_return** subroutines in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding. When the request is queued, the error status for the operation is **EINPROGRESS**. When all data has been successfully transferred, the error status is reset to reflect the success or failure of the operation. If the operation does not complete successfully, the error status for the operation is set to indicate the error. The *aio\_sigevent* member determines the asynchronous notification to occur when all operations have achieved synchronized I/O completion. All other members of the structure referenced by the *aiocbp* parameter are ignored. If the control block referenced by *aiocbp* becomes an illegal address prior to asynchronous I/O completion, the behavior is undefined.

If the **aio\_fsync** subroutine fails or *aiocbp* indicates an error condition, data is not guaranteed to have been successfully transferred.

## Parameters

<i>op</i>	Determines the way all currently queued I/O operations are completed.
<i>aiocbp</i>	Points to the <b>aiocb</b> structure associated with the I/O operation.

## aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

int	<i>aio_fildes</i>
off_t	<i>aio_offset</i>
char	* <i>aio_buf</i>
size_t	<i>aio_nbytes</i>
int	<i>aio_reqprio</i>
struct sigevent	<i>aio_sigevent</i>
int	<i>aio_lio_opcode</i>

## Execution Environment

The **aio\_error** and **aio\_error64** subroutines can be called from the process environment only.

## Return Values

The **aio\_fsync** subroutine returns a 0 to the calling process if the I/O operation is successfully queued. Otherwise, it returns a -1, and sets the **errno** global variable to indicate the error.

## Error Codes

<b>EAGAIN</b>	The requested asynchronous operation was not queued due to temporary resource limitations.
<b>EBADF</b>	The <i>aio_fildes</i> member of the <b>aiocb</b> structure referenced by the <i>aiocbp</i> parameter is not a valid file descriptor open for writing.

In the event that any of the queued I/O operations fail, the **aio\_fsync** subroutine returns the error condition defined for the **read** and **write** subroutines. The error is returned in the error status for the asynchronous **fsync** subroutine, which can be retrieved using the **aio\_error** subroutine.

## Related Information

“fcntl, dup, or dup2 Subroutine” on page 255, “fsync or fsync\_range Subroutine” on page 321, and “open, openx, open64, creat, or creat64 Subroutine” on page 935.

read, readx, readv, readvx, or pread Subroutine and write, writex, writev, writevx or pwrite Subroutines in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## aio\_nwait Subroutine

### Purpose

Suspends the calling process until a certain number of asynchronous I/O requests are completed.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <aio.h>

int aio_nwait (cnt, nwait, list)
int cnt;
int nwait;
struct aiocb **list;
```

### Description

Although the **aio\_nwait** subroutine is included with POSIX AIO, it is not part of the standard definitions for POSIX AIO.

The **aio\_nwait** subroutine suspends the calling process until a certain number (*nwait*) of asynchronous I/O requests are completed. These requests are initiated at an earlier time by the **lio\_listio** subroutine, which uses the LIO\_NOWAIT\_AIOWAIT *cmd* parameter. The **aio\_nwait** subroutine fills in the **aiocb** pointers to the completed requests in *list* and returns the number of valid entries in *list*. The *cnt* parameter is the maximum number of **aiocb** pointers that *list* can hold (*cnt* >= *nwait*). The subroutine also returns when less than *nwait* number of requests are done if there are no more pending aio requests.

**Note:** If the **lio\_listio64** subroutine is used, the **aiocb** structure changes to **aiocb64**.

**Note:** The aio control block's **errno** field continues to have the value EINPROG until after the **aio\_nwait** subroutine is completed. The **aio\_nwait** subroutine updates this field when the **lio\_listio** subroutine has run with the LIO\_NOWAIT\_AIOWAIT *cmd* parameter. No utility, such as **aio\_error**, can be used to look at this value until after the **aio\_nwait** subroutine is completed.

The **aio\_suspend** subroutine returns after any one of the specified requests gets done. The **aio\_nwait** subroutine returns after a certain number (*nwait* or more) of requests are completed.

There are certain limitations associated with the **aio\_nwait** subroutine, and a comparison between it and the **aio\_suspend** subroutine is necessary. The following table is a comparison of the two subroutines:

#### **aio\_suspend:**

Requires users to build a list of control blocks, each associated with an I/O operation they want to wait for.

Returns when any one of the specified control blocks indicates that the I/O associated with that control block completed.

#### **aio\_nwait:**

Requires the user to provide an array to put **aiocb** address information into. No specific aio control blocks need to be known.

Returns when *nwait* amount of requests are done or no other requests are to be processed.

**aio\_suspend:**

The aio control blocks may be updated before the subroutine is called. Other polling methods (such as the **aio\_error** subroutine) can also be used to view the aio control blocks.

**aio\_nwait:**

Updates the aio control blocks itself when it is called. Other polling methods can't be used until after the **aio\_nwait** subroutine is called enough times to cover all of the aio requests specified with the **lio\_listio** subroutine.

Is only used in accordance with the **LIO\_NOWAIT\_AIOWAIT** command, which is one of the commands associated with the **lio\_listio** subroutine. If the **lio\_listio** subroutine is not first used with the **LIO\_NOWAIT\_AIOWAIT** command, **aio\_nwait** can not be called. The **aio\_nwait** subroutine only affects those requests called by one or more **lio\_listio** calls for a specified process.

## Parameters

<i>cnt</i>	Specifies the number of entries in the list array. This number must be greater than 0 and less than 64 000.
<i>nwait</i>	Specifies the minimal number of requests to wait on. This number must be greater than 0 and less than or equal to the value specified by the <i>cnt</i> parameter.
<i>list</i>	An array of pointers to aio control structures defined in the <b>aio.h</b> file.

## Return Values

The return value is the total number of requests the **aio\_nwait** subroutine has waited on to complete. It can not be more than *cnt*. Although *nwait* is the desired amount of requests to find, the actual amount returned could be less than, equal to, or greater than *nwait*. The return value indicates how much of the list array to access.

The return value may be greater than the *nwait* value if the **lio\_listio** subroutine initiated more than *nwait* requests and the *cnt* variable is larger than *nwait*. The *nwait* parameter represents a minimal value desired for the return value, and *cnt* is the maximum value possible for the return.

The return value may be less than the *nwait* value if some of the requests initiated by the **lio\_listio** subroutine occur at a time of high activity, and there is a lack of resources available for the number of requests. **EAGAIN** (error try again later) may be returned in some request's aio control blocks, but these requests will not be seen by the **aio\_nwait** subroutine. In this situation **aioacb** addresses not found on the list have to be accessed by using the **aio\_error** subroutine after the **aio\_nwait** subroutine is called. You may need to increase the aio parameters *max servers* or *max requests* if this occurs. Increasing the parameters will ensure that the system is well tuned, and an **EAGAIN** error is less likely to occur.

In the event of an error, the **aio\_nwait** subroutine returns a value of -1 and sets the **errno** global variable to identify the error. Return codes can be set to the following **errno** values:

<b>EBUSY</b>	An <b>aio_nwait</b> call is in process.
<b>EINVAL</b>	The application has retrieved all of the <b>aioacb</b> pointers, but the user buffer does not have enough space for them.
<b>EINVAL</b>	There are no outstanding async I/O calls.
<b>EINVAL</b>	Specifies <i>cnt</i> or <i>nwait</i> values that are not valid.

## Related Information

The “aio\_cancel or aio\_cancel64 Subroutine” on page 38, “aio\_error or aio\_error64 Subroutine” on page 42, “aio\_nwait\_timeout Subroutine” on page 49, “aio\_read or aio\_read64 Subroutine” on page 50, “aio\_return or aio\_return64 Subroutine” on page 55, “aio\_suspend or aio\_suspend64 Subroutine” on page 58, “aio\_write or aio\_write64 Subroutine” on page 61, and “lio\_listio or lio\_listio64 Subroutine” on page 720.



The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## **aio\_nwait\_timeout Subroutine**

### **Purpose**

Extends the capabilities of the **aio\_nwait** subroutine by specifying timeout values.

### **Library**

Standard C library (**libc.a**).

### **Syntax**

```
int aio_nwait_timeout (cnt, nwait, list, timeout)
int cnt;
int nwait;
struct aiocbp **list;
int timeout;
```

### **Description**

The **aio\_nwait\_timeout** subroutine waits for a certain number of asynchronous I/O operations to complete as specified by the *nwait* parameter, or until the call has blocked for a certain duration specified by the *timeout* parameter.

### **Parameters**

<i>cnt</i>	Indicates the maximum number of pointers to the <b>aiocbp</b> structure that can be copied into the <b>list</b> array.
<i>list</i>	An array of pointers to aio control structures defined in the <b>aio.h</b> file.
<i>nwait</i>	Specifies the number of asynchronous I/O operations that must complete before the <b>aio_nwait_timeout</b> subroutine returns.
<i>timeout</i>	Specified in units of milliseconds.

A *timeout* value of -1 indicates that the subroutine behaves like the **aio\_nwait** subroutine, blocking until all of the requested I/O operations complete or until there are no more asynchronous I/O requests pending from the process.

A *timeout* value of 0 indicates that the subroutine returns immediately with the current completed number of asynchronous I/O requests. All other positive *timeout* values indicate that the subroutine must block until either the *timeout* value is reached or the requested number of asynchronous I/O operations complete.

### **Return Values**

The return value is the total number of requests the **aio\_nwait** subroutine has waited on to complete. It can not be more than *cnt*. Although *nwait* is the desired amount of requests to find, the actual amount returned could be less than, equal to, or greater than *nwait*. The return value indicates how much of the list array to access.

The return value may be greater than the *nwait* value if the **lio\_listio** subroutine initiated more than *nwait* requests and the *cnt* variable is larger than *nwait*. The *nwait* parameter represents a minimal value desired for the return value, and *cnt* is the maximum value possible for the return.

The return value may be less than the *nwait* value if some of the requests initiated by the **lio\_listio** subroutine occur at a time of high activity, and there is a lack of resources available for the number of requests. The **EAGAIN** return code (error try again later) might be returned in some request's aio control blocks, but these requests will not be seen by the **aio\_nwait** subroutine. In this situation, the **aioacb** structure addresses that are not found on the list must be accessed using the **aio\_error** subroutine after the **aio\_nwait** subroutine is called. You might need to increase the aio parameters max servers or max requests if this occurs. Increasing the parameters will ensure that the system is well tuned, and an **EAGAIN** error is less likely to occur. The return value might be less than the *nwait* value due to the setting of the new timeout parameter in the following cases:

- *timeout* > 0 and a timeout has occurred before *nwait* requests are done
- *timeout* = 0 and the current requests completed at the time of the **aio\_nwait\_timeout** call are less than *nwait* parameter

In the event of an error, the **aio\_nwait** subroutine returns a value of -1 and sets the **errno** global variable to identify the error. Return codes can be set to the following **errno** values:

- EBUSY** An **aio\_nwait** call is in process.
- EINVAL** The application has retrieved all of the **aioacb** pointers, but the user buffer does not have enough space for them.
- EINVAL** There are no outstanding async I/O calls.

## Related Information

“**aio\_nwait** Subroutine” on page 47, “**aio\_suspend** or **aio\_suspend64** Subroutine” on page 58, “**aio\_cancel** or **aio\_cancel64** Subroutine” on page 38, “**aio\_error** or **aio\_error64** Subroutine” on page 42, “**aio\_read** or **aio\_read64** Subroutine,” “**aio\_return** or **aio\_return64** Subroutine” on page 55, “**aio\_write** or **aio\_write64** Subroutine” on page 61, and “**lio\_listio** or **lio\_listio64** Subroutine” on page 720.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## **aio\_read** or **aio\_read64** Subroutine

The **aio\_read** or **aio\_read64** subroutine includes information for the POSIX AIO **aio\_read** subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO **aio\_read** subroutine.

### POSIX AIO **aio\_read** Subroutine

#### Purpose

Asynchronously reads a file.

#### Library

Standard C Library (**libc.a**)

#### Syntax

```
#include <aio.h>

int aio_read (aioacb)
struct aioacb *aioacb;
```

## Description

The **aio\_read** subroutine reads *aio\_nbytes* from the file associated with *aio\_fildes* into the buffer pointed to by *aio\_buf*. The subroutine returns when the read request has been initiated or queued to the file or device (even when the data cannot be delivered immediately).

The *aiocbp* value may be used as an argument to the **aio\_error** and **aio\_return** subroutines in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding. If an error condition is encountered during queuing, the function call returns without having initiated or queued the request. The requested operation takes place at the absolute position in the file as given by *aio\_offset*, as if the **lseek** subroutine were called immediately prior to the operation with an offset equal to *aio\_offset* and a whence equal to **SEEK\_SET**. After a successful call to enqueue an asynchronous I/O operation, the value of the file offset for the file is unspecified.

The *aio\_lio\_opcode* field is ignored by the **aio\_read** subroutine.

If prioritized I/O is supported for this file, the asynchronous operation is submitted at a priority equal to the scheduling priority of the process minus *aiocbp->aio\_reqprio*.

The *aiocbp* parameter points to an **aiocb** structure. If the buffer pointed to by *aio\_buf* or the control block pointed to by *aiocbp* becomes an illegal address prior to asynchronous I/O completion, the behavior is undefined.

Simultaneous asynchronous operations using the same *aiocbp* produce undefined results.

If synchronized I/O is enabled on the file associated with *aio\_fildes*, the behavior of this subroutine is according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion.

For any system action that changes the process memory space while an asynchronous I/O is outstanding, the result of that action is undefined.

For regular files, no data transfer occurs past the offset maximum established in the open file description.

If you use the **aio\_read** or **aio\_read64** subroutine with a file descriptor obtained from a call to the **shm\_open** subroutine, it will fail with **EINVAL**.

## Parameters

*aiocbp* Points to the **aiocb** structure associated with the I/O operation.

### aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
int          aio_fildes
off_t       aio_offset
char        *aio_buf
size_t      aio_nbytes
int         aio_reqprio
struct sigevent aio_sigevent
int         aio_lio_opcode
```

## Execution Environment

The **aio\_read** and **aio\_read64** subroutines can be called from the process environment only.

## Return Values

The `aio_read` subroutine returns 0 to the calling process if the I/O operation is successfully queued. Otherwise, it returns a -1 and sets the `errno` global variable to indicate the error.

## Error Codes

**EAGAIN**           The requested asynchronous I/O operation was not queued due to system resource limitations.

Each of the following conditions may be detected synchronously at the time of the call to the `aio_read` subroutine, or asynchronously. If any of the conditions below are detected synchronously, the `aio_read` subroutine returns -1 and sets the `errno` global variable to the corresponding value. If any of the conditions below are detected asynchronously, the return status of the asynchronous operation is set to -1, and the error status of the asynchronous operation is set to the corresponding value.

**EBADF**            The `aio_fildes` parameter is not a valid file descriptor open for reading.  
**EINVAL**            The file offset value implied by `aio_offset` is invalid, `aio_reqprio` is an invalid value, or `aio_nbytes` is an invalid value. The `aio_read` or `aio_read64` subroutine was used with a file descriptor obtained from a call to the `shm_open` subroutine.

If the `aio_read` subroutine successfully queues the I/O operation but the operation is subsequently canceled or encounters an error, the return status of the asynchronous operation is one of the values normally returned by the `read` subroutine. In addition, the error status of the asynchronous operation is set to one of the error statuses normally set by the `read` subroutine, or one of the following values:

**EBADF**            The `aio_fildes` argument is not a valid file descriptor open for reading.  
**ECANCELED**        The requested I/O was canceled before the I/O completed due to an explicit `aio_cancel` request.  
**EINVAL**            The file offset value implied by `aio_offset` is invalid.

The following condition may be detected synchronously or asynchronously:

**EOVERFLOW**        The file is a regular file, `aio_nbytes` is greater than 0, and the starting offset in `aio_offset` is before the end-of-file and is at or beyond the offset maximum in the open file description associated with `aio_fildes`.

## Related Information

“`aio_cancel` or `aio_cancel64` Subroutine” on page 38, “`aio_error` or `aio_error64` Subroutine” on page 42, “`aio_nwait` Subroutine” on page 47, “`aio_nwait_timeout` Subroutine” on page 49, “`lio_listio` or `lio_listio64` Subroutine” on page 720, “`aio_return` or `aio_return64` Subroutine” on page 55, “`aio_suspend` or `aio_suspend64` Subroutine” on page 58, “`aio_write` or `aio_write64` Subroutine” on page 61, “`close` Subroutine” on page 175, “`exec: execl, execlx, execlp, execv, execve, execvp, or exect` Subroutine” on page 236, “`exit, atexit, unatexit, _exit, or _Exit` Subroutine” on page 243, “`fork, f_fork, or vfork` Subroutine” on page 291, and “`lseek, llseek or lseek64` Subroutine” on page 765.

The `read`, `readx`, `readv`, `readvx`, or `pread` Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

### Legacy AIO `aio_read` Subroutine

## Purpose

Reads asynchronously from a file.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>
```

```
int aio_read( FileDescriptor, aiocbp)
int FileDescriptor;
struct aiocb *aiocbp;
```

```
int aio_read64( FileDescriptor, aiocbp)
int FileDescriptor;
struct aiocb64 *aiocbp;
```

## Description

The **aio\_read** subroutine reads asynchronously from a file. Specifically, the **aio\_read** subroutine reads from the file associated with the *FileDescriptor* parameter into a buffer.

The **aio\_read64** subroutine is similar to the **aio\_read** subroutine except that it takes an **aiocb64** reference parameter. This allows the **aio\_read64** subroutine to specify offsets in excess of **OFF\_MAX** (2 gigabytes minus 1).

In the large file enabled programming environment, **aio\_read** is redefined to be **aio\_read64** .

If you use the **aio\_read** or **aio\_read64** subroutine with a file descriptor obtained from a call to the **shm\_open** subroutine, it will fail with **EINVAL**.

The details of the read are provided by information in the **aiocb** structure, which is pointed to by the *aiocbp* parameter. This information includes the following fields:

<code>aio_buf</code>	Indicates the buffer to use.
<code>aio_nbytes</code>	Indicates the number of bytes to read.

When the read request has been queued, the **aio\_read** subroutine updates the file pointer specified by the `aio_whence` and `aio_offset` fields in the **aiocb** structure as if the requested I/O were already completed. It then returns to the calling program. The `aio_whence` and `aio_offset` fields have the same meaning as the *whence* and *offset* parameters in the **lseek** (“**lseek**, **llseek** or **lseek64** Subroutine” on page 765) subroutine. The subroutine ignores them for file objects that are not capable of seeking.

If an error occurs during the call, the read request is not queued. To determine the status of a request, use the **aio\_error** (“**aio\_error** or **aio\_error64** Subroutine” on page 42) subroutine.

To have the calling process receive the **SIGIO** signal when the I/O operation completes, set the **AIO\_SIGNAL** bit in the `aio_flag` field in the **aiocb** structure.

**Note:** The **event** structure in the **aiocb** structure is currently not in use but is included for future compatibility.

**Note:** The **\_AIO\_AIX\_SOURCE** macro used in **aio.h** must be defined when using **aio.h** to compile an aio application with the Legacy AIO function definitions. The default compile using the **aio.h** file is for an application with the POSIX AIO definitions. In the source file enter:

```
#define _AIO_AIX_SOURCE
#include <sys/aio.h>
```

or, on the command line when compiling enter:  
->xlc ... -D\_AIO\_AIX\_SOURCE ... legacy\_aio\_program.c

Since prioritized I/O is not supported at this time, the *aio\_reqprio* field of the structure is not presently used.

## Parameters

*FileDescriptor* Identifies the object to be read as returned from a call to open.  
*aiocbp* Points to the asynchronous I/O control block structure associated with the I/O operation.

## aiocb Structure

The **aiocb** and the **aiocb64** structures are defined in the **aio.h** file and contain the following members:

```
struct aiocb
{
    int          aio_whence;
    off_t       aio_offset;
    char        *aio_buf;
    ssize_t     aio_return;
    int         aio_errno;
    size_t      aio_nbytes;
    union {
        int      reqprio;
        struct {
            int  version:8;
            int  priority:8;
            int  cache_hint:16;
        } ext;
    } aio_ul;
    int         aio_flag;
    int         aio_iocpfd;
    aio_handle_t aio_handle;
}

#define aio_reqprio      aio_ul.reqprio
#define aio_version      aio_ul.ext.version
#define aio_priority     aio_ul.ext.priority
#define aio_cache_hint  aio_ul.ext.cache_hint
```

## Execution Environment

The **aio\_read** and **aio\_read64** subroutines can be called from the process environment only.

## Return Values

When the read request queues successfully, the **aio\_read** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the global variable **errno** to identify the error.

Return codes can be set to the following **errno** values:

**EAGAIN** Indicates that the system resources required to queue the request are not available. Specifically, the transmit queue may be full, or the maximum number of opens may be reached.

**EBADF** Indicates that the *FileDescriptor* parameter is not valid.

**EFAULT** Indicates that the address specified by the *aiocbp* parameter is not valid.

**EINVAL** Indicates that the *aio\_whence* field does not have a valid value, or that the resulting pointer is not valid. The **aio\_read** or **aio\_read64** subroutine was used with a file descriptor obtained from a call to the **shm\_open** subroutine.

When using I/O Completion Ports with AIO Requests, return codes can also be set to the following `errno` values:

<b>EBADF</b>	Indicates that the <code>aio_iocpfd</code> field in the <code>aio</code> structure is not a valid I/O Completion Port file descriptor.
<b>EINVAL</b>	Indicates that an I/O Completion Port service failed when attempting to start the AIO Request.
<b>EPERM</b>	Indicates that I/O Completion Port services are not available.

**Note:** Other error codes defined in the `sys/errno.h` file can be returned by the `aio_error` subroutine if an error during the I/O operation is encountered.

## Related Information

“`aio_cancel` or `aio_cancel64` Subroutine” on page 38, “`aio_nwait` Subroutine” on page 47, “`aio_nwait_timeout` Subroutine” on page 49, “`aio_error` or `aio_error64` Subroutine” on page 42, “`aio_return` or `aio_return64` Subroutine,” “`aio_suspend` or `aio_suspend64` Subroutine” on page 58, “`aio_write` or `aio_write64` Subroutine” on page 61, “`lio_listio` or `lio_listio64` Subroutine” on page 720.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## `aio_return` or `aio_return64` Subroutine

The `aio_return` or `aio_return64` subroutine includes information for the POSIX AIO `aio_return` subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO `aio_return` subroutine.

### POSIX AIO `aio_return` Subroutine

## Purpose

Retrieves the return status of an asynchronous I/O operation.

## Library

Standard C Library (`libc.a`)

## Syntax

```
#include <aio.h>
```

```
size_t aio_return (aioiop);  
struct aioiop *aioiop;
```

## Description

The `aio_return` subroutine returns the return status associated with the `aioiop` structure. The return status for an asynchronous I/O operation is the value that would be returned by the corresponding `read`, `write`, or `fsync` subroutine call. If the error status for the operation is equal to `EINPROGRESS`, the return status for the operation is undefined. The `aio_return` subroutine can be called once to retrieve the return status of a given asynchronous operation. After that, if the same `aioiop` structure is used in a call to `aio_return` or `aio_error`, an error may be returned. When the `aioiop` structure referred to by `aioiop` is used to submit another asynchronous operation, the `aio_return` subroutine can be successfully used to retrieve the return status of that operation.

## Parameters

*aiocbp* Points to the **aiocb** structure associated with the I/O operation.

## aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
int          aio_fildes
off_t       aio_offset
char        *aio_buf
size_t      aio_nbytes
int         aio_reqprio
struct sigevent aio_sigevent
int         aio_lio_opcode
```

## Execution Environment

The **aio\_return** and **aio\_return64** subroutines can be called from the process environment only.

## Return Values

If the asynchronous I/O operation has completed, the return status (as described for the **read**, **write**, and **fsync** subroutines) is returned. If the asynchronous I/O operation has not yet completed, the results of the **aio\_return** subroutine are undefined.

## Error Codes

**EINVAL** The *aiocbp* parameter does not refer to an asynchronous operation whose return status has not yet been retrieved.

## Related Information

“**aio\_cancel** or **aio\_cancel64** Subroutine” on page 38, “**aio\_error** or **aio\_error64** Subroutine” on page 42, “**aio\_nwait** Subroutine” on page 47, “**aio\_nwait\_timeout** Subroutine” on page 49, “**aio\_read** or **aio\_read64** Subroutine” on page 50, “**aio\_suspend** or **aio\_suspend64** Subroutine” on page 58, “**aio\_write** or **aio\_write64** Subroutine” on page 61, “**close** Subroutine” on page 175, “**exec: execl, execl, execlp, execv, execve, execvp, or exect** Subroutine” on page 236, “**exit, atexit, unatexit, \_exit, or \_Exit** Subroutine” on page 243, “**fork, f\_fork, or vfork** Subroutine” on page 291, “**lio\_listio** or **lio\_listio64** Subroutine” on page 720, and “**lseek, llseek** or **lseek64** Subroutine” on page 765.

The **read**, **readx**, **readv**, **readvx**, or **pread** Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

## Legacy AIO **aio\_return** Subroutine

## Purpose

Retrieves the return status of an asynchronous I/O request.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>
```

```
int aio_return( handle)
aio_handle_t handle;
```



```
int aio_return64( handle)
aio_handle_t handle;
```

## Description

The **aio\_return** subroutine retrieves the return status of the asynchronous I/O request associated with the **aio\_handle\_t** handle if the I/O request has completed. The status returned is the same as the status that would be returned by the corresponding **read** or **write** function calls. If the I/O operation has not completed, the returned status is undefined.

The **aio\_return64** subroutine is similar to the **aio\_return** subroutine except that it retrieves the error status associated with an **aiocb64** control block.

**Note:** The **\_AIO\_AIX\_SOURCE** macro used in **aio.h** must be defined when using **aio.h** to compile an aio application with the Legacy AIO function definitions. The default compile using the **aio.h** file is for an application with the POSIX AIO definitions. In the source file enter:

```
#define _AIO_AIX_SOURCE
#include <sys/aio.h>
```

or, on the command line when compiling enter:

```
->xlc ... -D_AIO_AIX_SOURCE ... legacy_aio_program.c
```

## Parameters

*handle*      The handle field of an **aio control block** (**aiocb** or **aiocb64**) structure is set by a previous call of the **aio\_read**, **aio\_read64**, **aio\_write**, **aio\_write64**, **lio\_listio**, **aio\_listio64** subroutine. If a random memory location is passed in, random results are returned.

## aiocb Structure

The **aiocb** structure is defined in the **/usr/include/aio.h** file and contains the following members:

```
struct aiocb
{
    int             aio_whence;
    off_t          aio_offset;
    char           *aio_buf;
    ssize_t        aio_return;
    int            aio_errno;
    size_t         aio_nbytes;
    union {
        int         reqprio;
        struct {
            int     version:8;
            int     priority:8;
            int     cache_hint:16;
        } ext;
    } aio_ul;
    int            aio_flag;
    int            aio_iocpfd;
    aio_handle_t  aio_handle;
}

#define aio_reqprio      aio_ul.reqprio
#define aio_version     aio_ul.ext.version
#define aio_priority    aio_ul.ext.priority
#define aio_cache_hint  aio_ul.ext.cache_hint
```

## Execution Environment

The **aio\_return** and **aio\_return64** subroutines can be called from the process environment only.

## Return Values

The `aio_return` subroutine returns the status of an asynchronous I/O request corresponding to those returned by `read` or `write` functions. If the error status returned by the `aio_error` subroutine call is `EINPROG`, the value returned by the `aio_return` subroutine is undefined.

## Examples

An `aio_read` request to read 1000 bytes from a disk device eventually, when the `aio_error` subroutine returns a 0, causes the `aio_return` subroutine to return 1000. An `aio_read` request to read 1000 bytes from a 500 byte file eventually causes the `aio_return` subroutine to return 500. An `aio_write` request to write to a read-only file system results in the `aio_error` subroutine eventually returning `EROFS` and the `aio_return` subroutine returning a value of -1.

## Related Information

“`aio_cancel` or `aio_cancel64` Subroutine” on page 38, “`aio_error` or `aio_error64` Subroutine” on page 42, “`aio_nwait` Subroutine” on page 47, “`aio_nwait_timeout` Subroutine” on page 49, “`aio_read` or `aio_read64` Subroutine” on page 50, “`aio_suspend` or `aio_suspend64` Subroutine,” “`aio_write` or `aio_write64` Subroutine” on page 61, “`close` Subroutine” on page 175, “`exec: execl, execl, execlp, execv, execve, execvp, or exect` Subroutine” on page 236, “`exit, atexit, unatexit, _exit, or _Exit` Subroutine” on page 243, “`fork, f_fork, or vfork` Subroutine” on page 291, “`lio_listio` or `lio_listio64` Subroutine” on page 720, and “`lseek, llseek` or `lseek64` Subroutine” on page 765.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## `aio_suspend` or `aio_suspend64` Subroutine

The `aio_suspend` subroutine includes information for the POSIX AIO `aio_suspend` subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO `aio_suspend` subroutine.

### POSIX AIO `aio_suspend` Subroutine

## Purpose

Waits for an asynchronous I/O request.

## Library

Standard C Library (`libc.a`)

## Syntax

```
#include <aio.h>

int aio_suspend (list, nent,
                timeout)
const struct aiocb * const list[];
int nent;
const struct timespec *timeout;
```

## Description

The `aio_suspend` subroutine suspends the calling thread until at least one of the asynchronous I/O operations referenced by the `list` parameter has completed, until a signal interrupts the function, or, if

timeout is not NULL, until the time interval specified by *timeout* has passed. If any of the **aio** structures in the list correspond to completed asynchronous I/O operations (the error status for the operation is not equal to **EINPROGRESS**) at the time of the call, the subroutine returns without suspending the calling thread. The *list* parameter is an array of pointers to asynchronous I/O control blocks. The *nent* parameter indicates the number of elements in the array. Each **aio** structure pointed to has been used in initiating an asynchronous I/O request through the **aio\_read**, **aio\_write**, or **lio\_listio** subroutine. This array may contain NULL pointers, which are ignored. If this array contains pointers that refer to **aio** structures that have not been used in submitting asynchronous I/O, the effect is undefined.

If the time interval indicated in the **timespec** structure pointed to by *timeout* passes before any of the I/O operations referenced by *list* are completed, the **aio\_suspend** subroutine returns with an error. If the Monotonic Clock option is supported, the clock that is used to measure this time interval is the **CLOCK\_MONOTONIC** clock.

## Parameters

<i>list</i>	Array of asynchronous I/O operations.
<i>nent</i>	Indicates the number of elements in the <i>list</i> array.
<i>timeout</i>	Specifies the time the subroutine has to complete the operation.

## Execution Environment

The **aio\_suspend** and **aio\_suspend64** subroutines can be called from the process environment only.

## Return Values

If the **aio\_suspend** subroutine returns after one or more asynchronous I/O operations have completed, it returns a 0. Otherwise, it returns a -1 and sets the **errno** global variable to indicate the error.

The application can determine which asynchronous I/O completed by scanning the associated error and returning status using the **aio\_error** and **aio\_return** subroutines, respectively.

## Error Codes

<b>EAGAIN</b>	No asynchronous I/O indicated in the list referenced by <i>list</i> completed in the time interval indicated by <i>timeout</i> .
<b>EINTR</b>	A signal interrupted the <b>aio_suspend</b> subroutine. Since each asynchronous I/O operation may possibly provoke a signal when it completes, this error return may be caused by the completion of one (or more) of the very I/O operations being awaited.

## Related Information

“**aio\_cancel** or **aio\_cancel64** Subroutine” on page 38, “**aio\_error** or **aio\_error64** Subroutine” on page 42, “**aio\_nwait** Subroutine” on page 47, “**aio\_nwait\_timeout** Subroutine” on page 49, “**aio\_read** or **aio\_read64** Subroutine” on page 50, “**aio\_return** or **aio\_return64** Subroutine” on page 55, “**aio\_write** or **aio\_write64** Subroutine” on page 61, and “**lio\_listio** or **lio\_listio64** Subroutine” on page 720.

## Legacy AIO **aio\_suspend** Subroutine

## Purpose

Suspends the calling process until one or more asynchronous I/O requests is completed.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>
```

```
aio_suspend( count, aiocbpa)
int count;
struct aiocb *aiocbpa[ ];
```

```
aio_suspend64( count, aiocbpa)
int count;
struct aiocb64 *aiocbpa[ ];
```

## Description

The **aio\_suspend** subroutine suspends the calling process until one or more of the *count* parameter asynchronous I/O requests are completed or a signal interrupts the subroutine. Specifically, the **aio\_suspend** subroutine handles requests associated with the **aio control block (aiocb)** structures pointed to by the *aiocbpa* parameter.

The **aio\_suspend64** subroutine is similar to the **aio\_suspend** subroutine except that it takes an array of pointers to **aiocb64** structures. This allows the **aio\_suspend64** subroutine to suspend on asynchronous I/O requests submitted by either the **aio\_read64**, **aio\_write64**, or the **lio\_listio64** subroutines.

In the large file enabled programming environment, **aio\_suspend** is redefined to be **aio\_suspend64**.

The array of **aiocb** pointers may include null pointers, which will be ignored. If one of the I/O requests is already completed at the time of the **aio\_suspend** call, the call immediately returns.

**Note:** The `_AIO_AIX_SOURCE` macro used in **aio.h** must be defined when using **aio.h** to compile an aio application with the Legacy AIO function definitions. The default compile using the **aio.h** file is for an application with the POSIX AIO definitions. In the source file enter:

```
#define _AIO_AIX_SOURCE
#include <sys/aio.h>
```

or, on the command line when compiling enter:

```
->xlc ... -D_AIO_AIX_SOURCE ... legacy_aio_program.c
```

## Parameters

*count* Specifies the number of entries in the *aiocbpa* array.

*aiocbpa* Points to the **aiocb** or **aiocb64** structures associated with the asynchronous I/O operations.

## aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
struct aiocb
{
    int          aio_whence;
    off_t       aio_offset;
    char        *aio_buf;
    ssize_t     aio_return;
    int         aio_errno;
    size_t      aio_nbytes;
    union {
        int      reqprio;
        struct {
            int  version:8;
            int  priority:8;
            int  cache_hint:16;
        };
    };
};
```

```

        } ext;
    } aio_ul;
    int         aio_flag;
    int         aio_iocpfd;
    aio_handle_t aio_handle;
}

#define aio_reqprio    aio_ul.reqprio
#define aio_version    aio_ul.ext.version
#define aio_priority   aio_ul.ext.priority
#define aio_cache_hint aio_ul.ext.cache_hint

```

## Execution Environment

The **aio\_suspend** and **aio\_suspend64** subroutines can be called from the process environment only.

## Return Values

If one or more of the I/O requests completes, the **aio\_suspend** subroutine returns the index into the `aio_cbpa` array of one of the completed requests. The index of the first element in the `aio_cbpa` array is 0. If more than one request has completed, the return value can be the index of any of the completed requests.

In the event of an error, the **aio\_suspend** subroutine returns a value of -1 and sets the **errno** global variable to identify the error. Return codes can be set to the following **errno** values:

**EINTR** Indicates that a signal or event interrupted the **aio\_suspend** subroutine call.  
**EINVAL** Indicates that the `aio_whence` field does not have a valid value or that the resulting pointer is not valid.

## Related Information

“`aio_cancel` or `aio_cancel64` Subroutine” on page 38, “`aio_error` or `aio_error64` Subroutine” on page 42, “`aio_nwait` Subroutine” on page 47, “`aio_nwait_timeout` Subroutine” on page 49, “`aio_read` or `aio_read64` Subroutine” on page 50, “`aio_return` or `aio_return64` Subroutine” on page 55, “`aio_write` or `aio_write64` Subroutine,” and “`lio_listio` or `lio_listio64` Subroutine” on page 720.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## **aio\_write** or **aio\_write64** Subroutine

The **aio\_write** subroutine includes information for the POSIX AIO **aio\_write** subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO **aio\_write** subroutine.

### POSIX AIO **aio\_write** Subroutine

#### Purpose

Asynchronously writes to a file.

#### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <aio.h>
```

```
int aio_write (aiocbp)  
struct aiocb *aiocbp;
```

## Description

The **aio\_write** subroutine writes *aio\_nbytes* to the file associated with *aio\_fildes* from the buffer pointed to by *aio\_buf*. The subroutine returns when the write request has been initiated or queued to the file or device.

The *aiocbp* parameter may be used as an argument to the **aio\_error** and **aio\_return** subroutines in order to determine the error status and return status, respectively, of the asynchronous operation while it is proceeding.

The *aiocbp* parameter points to an **aiocb** structure. If the buffer pointed to by *aio\_buf* or the control block pointed to by *aiocbp* becomes an illegal address prior to asynchronous I/O completion, the behavior is undefined.

If **O\_APPEND** is not set for the *aio\_fildes* file descriptor, the requested operation takes place at the absolute position in the file as given by *aio\_offset*. This is done as if the **lseek** subroutine were called immediately prior to the operation with an offset equal to *aio\_offset* and a whence equal to **SEEK\_SET**. If **O\_APPEND** is set for the file descriptor, write operations append to the file in the same order as the calls were made. After a successful call to enqueue an asynchronous I/O operation, the value of the file offset for the file is unspecified.

The *aio\_lio\_opcode* field is ignored by the **aio\_write** subroutine.

If prioritized I/O is supported for this file, the asynchronous operation is submitted at a priority equal to the scheduling priority of the process minus *aiocbp->aio\_reqprio*.

Simultaneous asynchronous operations using the same *aiocbp* produce undefined results.

If synchronized I/O is enabled on the file associated with *aio\_fildes*, the behavior of this subroutine is according to the definitions of synchronized I/O data integrity completion, and synchronized I/O file integrity completion.

For any system action that changes the process memory space while an asynchronous I/O is outstanding, the result of that action is undefined.

For regular files, no data transfer occurs past the offset maximum established in the open file description associated with *aio\_fildes*.

If you use the **aio\_write** or **aio\_write64** subroutine with a file descriptor obtained from a call to the **shm\_open** subroutine, it will fail with **EINVAL**.

## Parameters

*aiocbp* Points to the **aiocb** structure associated with the I/O operation.

### aiocb Structure

The **aiocb** structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
int          aio_fildes  
off_t       aio_offset  
char        *aio_buf
```

size_t	aio_nbytes
int	aio_reqprio
struct sigevent	aio_sigevent
int	aio_lio_opcode

## Execution Environment

The **aio\_write** and **aio\_write64** subroutines can be called from the process environment only.

## Return Values

The **aio\_write** subroutine returns a 0 to the calling process if the I/O operation is successfully queued. Otherwise, a -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

**EAGAIN** The requested asynchronous I/O operation was not queued due to system resource limitations.

Each of the following conditions may be detected synchronously at the time of the call to **aio\_write**, or asynchronously. If any of the conditions below are detected synchronously, the **aio\_write** subroutine returns a -1 and sets the **errno** global variable to the corresponding value. If any of the conditions below are detected asynchronously, the return status of the asynchronous operation is set to -1, and the error status of the asynchronous operation is set to the corresponding value.

**EBADF** The *aio\_fildes* parameter is not a valid file descriptor open for writing.  
**EINVAL** The file offset value implied by *aio\_offset* is invalid, *aio\_reqprio* is an invalid value, or *aio\_nbytes* is an invalid value. The **aio\_write** or **aio\_write64** subroutine was used with a file descriptor obtained from a call to the **shm\_open** subroutine.

If the **aio\_write** subroutine successfully queues the I/O operation, the return status of the asynchronous operation is one of the values normally returned by the **write** subroutine call. If the operation is successfully queued but is subsequently canceled or encounters an error, the error status for the asynchronous operation contains one of the values normally set by the **write** subroutine call, or one of the following:

**EBADF** The *aio\_fildes* parameter is not a valid file descriptor open for writing.  
**EINVAL** The file offset value implied by *aio\_offset* would be invalid.  
**ECANCELED** The requested I/O was canceled before the I/O completed due to an **aio\_cancel** request.

The following condition may be detected synchronously or asynchronously:

**EFBIG** The file is a regular file, *aio\_nbytes* is greater than 0, and the starting offset in *aio\_offset* is at or beyond the offset maximum in the open file description associated with *aio\_fildes*.

## Related Information

“aio\_cancel or aio\_cancel64 Subroutine” on page 38, “aio\_error or aio\_error64 Subroutine” on page 42, “aio\_nwait Subroutine” on page 47, “aio\_nwait\_timeout Subroutine” on page 49, “lio\_listio or lio\_listio64 Subroutine” on page 720, “aio\_read or aio\_read64 Subroutine” on page 50, “aio\_suspend or aio\_suspend64 Subroutine” on page 58, “aio\_return or aio\_return64 Subroutine” on page 55, “close Subroutine” on page 175, “exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236, “exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243, “fork, f\_fork, or vfork Subroutine” on page 291, and “lseek, llseek or lseek64 Subroutine” on page 765.

The read, readx, readv, readvx, or pread Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

## Legacy AIO `aio_write` Subroutine

### Purpose

Writes to a file asynchronously.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <aio.h>
```

```
int aio_write( FileDescriptor, aiocbp)
int FileDescriptor;
struct aiocb *aiocbp;
```

```
int aio_write64( FileDescriptor, aiocbp)
int FileDescriptor;
struct aiocb64 *aiocbp;
```

### Description

The **aio\_write** subroutine writes asynchronously to a file. Specifically, the **aio\_write** subroutine writes to the file associated with the *FileDescriptor* parameter from a buffer. To handle this, the subroutine uses information from the **aio control block (aiocb)** structure, which is pointed to by the *aiocbp* parameter. This information includes the following fields:

<code>aio_buf</code>	Indicates the buffer to use.
<code>aio_nbytes</code>	Indicates the number of bytes to write.

The **aio\_write64** subroutine is similar to the **aio\_write** subroutine except that it takes an **aiocb64** reference parameter. This allows the **aio\_write64** subroutine to specify offsets in excess of `OFF_MAX` (2 gigabytes minus 1).

In the large file enabled programming environment, **aio\_read** is redefined to be **aio\_read64**.

If you use the **aio\_write** or **aio\_write64** subroutine with a file descriptor obtained from a call to the **shm\_open** subroutine, it will fail with **EINVAL**.

When the write request has been queued, the **aio\_write** subroutine updates the file pointer specified by the `aio_whence` and `aio_offset` fields in the **aiocb** structure as if the requested I/O completed. It then returns to the calling program. The `aio_whence` and `aio_offset` fields have the same meaning as the *whence* and *offset* parameters in the **lseek** (“*lseek*, *llseek* or *lseek64* Subroutine” on page 765) subroutine. The subroutine ignores them for file objects that are not capable of seeking.

If an error occurs during the call, the write request is not initiated or queued. To determine the status of a request, use the **aio\_error** (“*aio\_error* or *aio\_error64* Subroutine” on page 42) subroutine.

To have the calling process receive the **SIGIO** signal when the I/O operation completes, set the `AIO_SIGNAL` bit in the `aio_flag` field in the **aiocb** structure.

**Note:** The **event** structure in the **aiocb** structure is currently not in use but is included for future compatibility.



**Note:** The `_AIO_AIX_SOURCE` macro used in `aio.h` must be defined when using `aio.h` to compile an aio application with the Legacy AIO function definitions. The default compile using the `aio.h` file is for an application with the POSIX AIO definitions. In the source file enter:

```
#define _AIO_AIX_SOURCE
#include <sys/aio.h>
```

or, on the command line when compiling enter:

```
->xlc ... -D_AIO_AIX_SOURCE ... legacy_aio_program.c
```

Since prioritized I/O is not supported at this time, the `aio_reqprio` field of the structure is not presently used.

## Parameters

*FileDescriptor*      Identifies the object to be written as returned from a call to open.  
*aiocbp*                Points to the asynchronous I/O control block structure associated with the I/O operation.

## aiocb Structure

The `aiocb` structure is defined in the `/usr/include/aio.h` file and contains the following members:

```
struct aiocb
{
    int             aio_whence;
    off_t           aio_offset;
    char            *aio_buf;
    ssize_t         aio_return;
    int             aio_errno;
    size_t          aio_nbytes;
    union {
        int         reqprio;
        struct {
            int     version:8;
            int     priority:8;
            int     cache_hint:16;
        } ext;
    } aio_ul;
    int             aio_flag;
    int             aio_iocpfd;
    aio_handle_t    aio_handle;
}

#define aio_reqprio      aio_ul.reqprio
#define aio_version      aio_ul.ext.version
#define aio_priority     aio_ul.ext.priority
#define aio_cache_hint  aio_ul.ext.cache_hint
```

## Execution Environment

The `aio_write` and `aio_write64` subroutines can be called from the process environment only.

## Return Values

When the write request queues successfully, the `aio_write` subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the `errno` global variable to identify the error.

Return codes can be set to the following `errno` values:

**EAGAIN**            Indicates that the system resources required to queue the request are not available. Specifically, the transmit queue may be full, or the maximum number of opens may have been reached.  
**EBADF**            Indicates that the *FileDescriptor* parameter is not valid.  
**EFAULT**           Indicates that the address specified by the *aiocbp* parameter is not valid.

**EINVAL** Indicates that the `aio_whence` field does not have a valid value or that the resulting pointer is not valid. The `aio_write` or `aio_write64` subroutine was used with a file descriptor obtained from a call to the `shm_open` subroutine.

When using I/O Completion Ports with AIO Requests, return codes can also be set to the following `errno` values:

**EBADF** Indicates that the `aio_iocpfd` field in the `aio_cb` structure is not a valid I/O Completion Port file descriptor.

**EINVAL** Indicates that an I/O Completion Port service failed when attempting to start the AIO Request.

**EPERM** Indicates that I/O Completion Port services are not available.

**Note:** Other error codes defined in the `/usr/include/sys/errno.h` file may be returned by the `aio_error` subroutine if an error during the I/O operation is encountered.

## Related Information

“`aio_cancel` or `aio_cancel64` Subroutine” on page 38, “`aio_error` or `aio_error64` Subroutine” on page 42, “`aio_nwait` Subroutine” on page 47, “`aio_nwait_timeout` Subroutine” on page 49, “`aio_read` or `aio_read64` Subroutine” on page 50, “`aio_return` or `aio_return64` Subroutine” on page 55, “`aio_suspend` or `aio_suspend64` Subroutine” on page 58, “`lio_listio` or `lio_listio64` Subroutine” on page 720.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## **alloc, dealloc, print, read\_data, read\_regs, symbol\_addr, write\_data, and write\_regs Subroutine**

### Purpose

Provide access to facilities needed by the `pthread debug` library and supplied by the debugger or application.

### Library

`pthread debug` library (`libpthdebug.a`)

### Syntax

```
#include <sys/pthdebug.h>

int alloc (user, len, bufp)
pthdb_user_t user;
size_t len;
void **bufp;

int dealloc (user, buf)
pthdb_user_t user;
void *buf;

int print (user, str)
pthdb_user_t user;
char *str;
```

```

int read_data (user, buf, addr, size)
pthdb_user_t user;
void *buf;
pthdb_addr_t addr;
int size;

int read_regs (user, tid, flags, context)
pthdb_user_t user;
tid_t tid;
unsigned long long flags;
struct context64 *context;

int symbol_addrs (user, symbols[], count)
pthdb_user_t user;
pthdb_symbol_t symbols[];
int count;

int write_data (user, buf, addr, size)
pthdb_user_t user;
void *buf;
pthdb_addr_t addr;
int size;

int write_regs (user, tid, flags, context)
pthdb_user_t user;
tid_t tid;
unsigned long long flags;
struct context64 *context;

```

## Description

### *int alloc()*

Allocates *len* bytes of memory and returns the address. If successful, 0 is returned; otherwise, a nonzero number is returned. This call back function is always required.

### *int dealloc()*

Takes a buffer and frees it. If successful, 0 is returned; otherwise, a nonzero number is returned. This call back function is always required.

### *int print()*

Prints the character string to the debugger's stdout. If successful, 0 is returned; otherwise, a nonzero number is returned. This call back is for debugging the library only. If you aren't debugging the pthread debug library, pass a NULL value for this call back.

### *int read\_data()*

Reads the requested number of bytes of data at the requested location from an active process or core file and returns the data through a buffer. If successful, 0 is returned; otherwise, a nonzero number is returned. This call back function is always required.

### *int read\_regs()*

Reads the context information of a debuggee kernel thread from an active process or from a core file. The information should be formatted in **context64** form for both a 32-bit and a 64-bit process. If successful, 0 is returned; otherwise, a nonzero number is returned. This function is only required when using the **pthdb\_thread\_context** and **pthdb\_thread\_setcontext** subroutines.

### *int symbol\_addrs()*

Resolves the address of symbols in the debuggee. The pthread debug library calls this subroutine to get the address of known debug symbols. If the symbol has a name of NULL or "", set the address to 0LL instead of doing a lookup or returning an error. If successful, 0 is returned; otherwise, a nonzero number is returned. In introspective mode, when the **PTHDB\_FLAG\_SUSPEND** flag is set, the application can use the pthread debug library by passing NULL, or it can use one of its own.

### *int write\_data()*

Writes the requested number of bytes of data to the requested location. The **libpthdebug.a** library may use this to write data to the active process. If successful, 0 is returned; otherwise, a nonzero

number is returned. This call back function is required when the **PTHDB\_FLAG\_HOLD** flag is set and when using the **pthdb\_thread\_setcontext** subroutine.

*int write\_regs()*

Writes requested context information to specified debuggee's kernel thread id. If successful, 0 is returned; otherwise, a nonzero number is returned. This subroutine is only required when using the **pthdb\_thread\_setcontext** subroutine.

**Note:** If the **write\_data** and **write\_regs** subroutines are NULL, the pthread debug library will not try to write data or regs. If the **pthdb\_thread\_set\_context** subroutine is called when the **write\_data** and **write\_regs** subroutines are NULL, **PTHDB\_NOTSUP** is returned.

## Parameters

<i>user</i>	User handle.
<i>symbols</i>	Array of symbols.
<i>count</i>	Number of symbols.
<i>buf</i>	Buffer.
<i>addr</i>	Address to be read from or wrote to.
<i>size</i>	Size of the buffer.
<i>flags</i>	Session flags, must accept <b>PTHDB_FLAG_GPRS</b> , <b>PTHDB_FLAG_SPRS</b> , <b>PTHDB_FLAG_FPRS</b> , and <b>PTHDB_FLAG_REGS</b> .
<i>tid</i>	Thread id.
<i>flags</i>	Flags that control which registers are read or wrote.
<i>context</i>	Context structure.
<i>len</i>	Length of buffer to be allocated or reallocated.
<i>bufp</i>	Pointer to buffer.
<i>str</i>	String to be printed.

## Return Values

If successful, these subroutines return 0; otherwise they return a nonzero value.

## Related Information

The “malloc, free, realloc, calloc, malloc, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778.

---

## alloclmb Subroutine

### Purpose

Allocates a contiguous block of contiguous real memory for exclusive use by the caller. The block of memory reserved will be the size of a system LMB.

### Syntax

```
#include <sys/dr.h>
```

```
int alloclmb(long long *laddr, int flags)
```

### Description

The **alloclmb()** subroutine reserves an LMB sized block of contiguous real memory for exclusive use by the caller. It returns the partition logical address of that memory in *\*laddr*.

**allocmb()** is only valid in an LPAR environment, and it fails (with **ENOTSUP**) if called in another environment.

Only a privileged user should call **allocmb()**.

## Parameters

*laddr*                On successful return, contains the logical address of the allocated LMB.  
*flags*                Must be 0.

## Execution Environment

This **allocmb()** interface should only be called from the process environment.

## Return Values

0                      The LMB is successfully allocated.

## Error Codes

**ENOTSUP**            LMB allocation not supported on this system.  
**EINVAL**             Invalid flags value.  
**EINVAL**             Not in the process environment.  
**ENOMEM**            A free LMB could not be made available.

## Related Information

“freelmb Subroutine” on page 313

---

## arm\_end Subroutine

### Purpose

The **arm\_end** subroutine is used to mark the end of an application. This subroutine call must always be called when a program that issued an **arm\_init** (“arm\_init Subroutine” on page 77) subroutine call terminates. In the PTX<sup>®</sup> implementation of ARM, application data structures may persist after **arm\_end** is issued.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h
```

```
arm_ret_stat_t ARM_API arm_end( arm_appl_id_t appl_id,        /* application id  
*/  
          arm_flag_t flags,                                /* Reserved = 0        */  
          arm_data_t *data,                               /* Reserved = NULL    */  
          arm_data_sz_t data_size);                      /* Reserved = 0        */
```

## Description

By calling the **arm\_end** subroutine, an application program signals to the ARM library that it has ceased issuing ARM subroutine calls for the application specified and that the library code can remove references to the application. As far as the calling program is concerned, all references to transactions defined for the named application can be removed as well.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

Note that, in the PTX implementation of ARM, multiple processes can issue **arm\_init** (“arm\_init Subroutine” on page 77) subroutine calls for a given application with the effect that multiple simultaneous definitions of the application are effective. The ARM library code points all these definitions to a single application structure in the ARM private shared memory area. A use-count keeps track of the number of simultaneous definitions. Each time **arm\_init** is issued for the application name, the counter is incremented and each time the **arm\_end** subroutine call is issued for the associated *appl\_id*, the counter is decremented. No call to **arm\_end** is permitted to decrement the counter less than zero.

Only when the counter reaches zero is the application structure inactivated. As long as the counter is non-zero, transactions defined for the application remain active and new transactions can be defined for the application. It does not matter which process created the definition of the application.

This implementation was chosen because it makes perfect sense in a PTX environment. Any more restrictive implementation would have increased memory use significantly and would be useless for PTX monitoring purposes.

## Parameters

### **appl\_id**

The identifier is returned by an earlier call to **arm\_init**, “arm\_init Subroutine” on page 77. The PTX implementation does not require that the **arm\_init** subroutine call was issued by the same program or process now issuing the **arm\_end** subroutine call. However, each time the **arm\_end** subroutine call is issued against an *appl\_id*, the use-count of the transaction structure is decremented. When the count reaches zero, the application structure and all associated transaction structures are marked as inactive. Subsequent **arm\_init** calls can reactivate the application structure but transaction structures formerly associated with the application are not automatically activated. Each transaction must be reactivated through the **arm\_getid** (“arm\_getid Subroutine” on page 73) subroutine call.

The *appl\_id* is used to look for an application structure. If none is found, no action is taken and the function returns -1. If one is found, the use-count of the application structure is decremented. If that makes the counter zero, the use-counts of all associated transaction structures are set to zero. The total number of application structures that have been initialized for the calling process but not ended is decremented. If this count reaches zero, access to the shared memory from the process is released and the count of users of the shared memory area is decremented. If the count of users of the shared memory segment reaches zero, the shared memory segment is deleted.

### **flags, data, data\_size**

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns zero. If the subroutine fails, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

`/usr/include/arm.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

1198, **arm\_init** (“arm\_init Subroutine” on page 77) subroutine, **arm\_getid** (“arm\_getid Subroutine” on page 73) subroutine.

---

## arm\_end Dual Call Subroutine

### Purpose

The **arm\_end** subroutine is used to mark the end of an application. This subroutine call must always be called when a program that issued an **arm\_init** (“arm\_init Dual Call Subroutine” on page 79) subroutine call terminates. In the PTX implementation of ARM, application data structures may persist after **arm\_end** is issued. This may not be the case for the *lower library* implementation.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_ret_stat_t ARM_API arm_end( arm_appl_id_t appl_id,      /* application id
*/
    arm_flag_t flags,          /* Reserved = 0          */
    arm_data_t *data,          /* Reserved = NULL      */
    arm_data_sz_t data_size);  /* Reserved = 0          */
```

### Description

By calling the **arm\_end** subroutine, an application program signals to the ARM library that it has ceased issuing ARM subroutine calls for the application specified and that the library code can remove references to the application. As far as the calling program is concerned, all references to transactions defined for the named application can be removed as well.

Before the PTX implementation code is executed, the *lower library* is called. If this call returns a value of zero, that return value is passed to the caller. If the value returned by the *lower library* is non-zero, the return value is the one generated by the PTX library code.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

Note that, in the PTX implementation of ARM, multiple processes can issue **arm\_init** (“arm\_init Dual Call Subroutine” on page 79) subroutine calls for a given application with the effect that multiple simultaneous definitions of the application are effective. The ARM library code points all these definitions to a single application structure in the ARM private shared memory area. A use-count keeps track of the number of simultaneous definitions. Each time **arm\_init** is issued for the application name, the counter is incremented and each time the **arm\_end** subroutine call is issued for the associated *appl\_id*, the counter is decremented. No call to **arm\_end** is permitted to decrement the counter less than zero.

Only when the counter reaches zero is the application structure inactivated. As long as the counter is non-zero, transactions defined for the application remain active and new transactions can be defined for the application. It does not matter which process created the definition of the application.

This implementation was chosen because it makes perfect sense in a PTX environment. Any more restrictive implementation would have increased memory use significantly and would be useless for PTX monitoring purposes.

For the implementation of **arm\_end** in the *lower library*, other restrictions may exist.

## Parameters

### appl\_id

The identifier returned by an earlier call to **arm\_init** (“arm\_init Dual Call Subroutine” on page 79). The identifier is passed to the **arm\_end** function of the *lower library*. If the *lower library* returns a zero, a zero is returned to the caller. After the invocation of the *lower library*, the PTX implementation attempts to translate the *appl\_id* argument to its own identifier from the cross-reference table created by **arm\_init** (“arm\_init Dual Call Subroutine” on page 79). If one can be found, it is used for the PTX implementation; if no cross reference is found, the *appl\_id* is used as passed in. The PTX implementation does not require that the **arm\_init** subroutine call was issued by the same program or process now issuing the **arm\_end** subroutine call. However, each time the **arm\_end** subroutine call is issued against an *appl\_id*, the use-count of the transaction structure is decremented. When the count reaches zero, the application structure and all associated transaction structures are marked as inactive. Subsequent **arm\_init** calls can reactivate the application structure but transaction structures formerly associated with the application are not automatically activated. Each transaction must be reactivated through the **arm\_getid** (“arm\_getid Dual Call Subroutine” on page 75) subroutine call.

In the PTX implementation, the *appl\_id* (as retrieved from the cross-reference table) is used to look for an application structure. If none is found, no action is taken and the PTX function is considered to have failed. If one is found, the use-count of the application structure is decremented. If that makes the counter zero, the use-counts of all associated transaction structures are set to zero. The total number of application structures that have been initialized for the calling process but not ended is decremented. If this count reaches zero, access to the shared memory from the process is released and the count of users of the shared memory area is decremented. If the count of users of the shared memory segment reaches zero, the shared memory segment is deleted.

### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns zero. If the subroutine fails, a value less than zero is returned. If the call to the *lower library* was successful, a zero is returned. If the subroutine call to the *lower library* failed but the PTX implementation didn't fail, a zero is returned. If both implementations failed, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

<b>/usr/include/arm.h</b>	Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.
---------------------------	---



## Related Information

- “arm\_init Dual Call Subroutine” on page 79
- “arm\_getid Dual Call Subroutine” on page 75

---

## arm\_getid Subroutine

### Purpose

The **arm\_getid** subroutine is used to register a transaction as belonging to an application and assign a unique identifier to the application/transaction pair. In the PTX implementation of ARM, multiple instances of a transaction within one application can't be defined. A transaction must be registered before any ARM measurements can begin.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_tran_id_t arm_getid(      arm_appl_id_t appl_id,      /* application handle
*/
    arm_ptr_t    *tran_name,    /* transaction name          */
    arm_ptr_t    *tran_detail,  /* transaction additional info */
    arm_flag_t   flags,         /* Reserved = 0              */
    arm_data_t   *data,         /* Reserved = NULL           */
    arm_data_sz_t data_size);  /* Reserved = 0              */
```

### Description

Each transaction needs to be defined by a unique name within an application. Transactions can be defined so they best fit the application environment. For example, if a given environment has thousands of unique transactions, it may be feasible to define groups of similar transactions to prevent data overload. In other situations, you may want to use generated transaction names that reflect what data a transaction carries along with the transaction type. For example, the type of SQL query could be analyzed to group customer query transactions according to complexity, such as *customer\_simple*, *customer*, *customer\_complex*.

Whichever method is used to name transactions, in the PTX implementation of the ARM API, measurements are always collected for each unique combination of:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

Note that the use-count for a transaction structure is either one or zero. This ensures that as long as the application structure is active, so are all transactions for which an **arm\_getid** call was issued after the application was activated by an **arm\_init** (“arm\_init Subroutine” on page 77) call. The transaction use-count is reset to zero by the **arm\_end** (“arm\_end Subroutine” on page 69) call if this call causes the application use-count to go to zero.

Note that the implementation of **arm\_getid** doesn't allow unique instances of a transaction to be defined. The **tran\_id** associated with a transaction is stored in the ARM shared memory area and will remain constant throughout the life of the shared memory area. Consequently, subsequent executions of a

program that defines one or more transactions under a given application will usually have the same ID returned for the transactions each time. The same is true when different programs define the same transaction within an application: As long as the shared memory area exists, they will all have the same ID returned. This is done to minimize the use of memory for transaction definitions and because it makes no difference from a PTX point of view.

If this is not acceptable from an application point of view, programs can dynamically generate transaction names to pass on the **arm\_getid** subroutine call.

## Parameters

### **appl\_id**

The identifier returned by an earlier call to **arm\_init** (“arm\_init Subroutine” on page 77). The PTX implementation does not require that the **arm\_init** subroutine call was issued by the same program or process now issuing the **arm\_getid** subroutine call. However, the number of issued **arm\_init** subroutine calls for the application name must exceed the number of issued **arm\_end** subroutine calls for this *appl\_id*.

The *appl\_id* is used to look for an application structure. If one is not found or if the use-count of the one found is zero, no action is taken and the function returns -1.

### **tran\_name**

A unique transaction name. The name only needs to be unique within the *appl\_id*. The maximum length is 128 characters including the terminating zero. The argument is converted to a key by removing all blanks and truncating the string to 32 characters, including a terminating zero. This key is used to look for a transaction structure (that belongs to the application identified in the first argument) in the library’s private shared memory area. If a transaction structure is found, its use-count is set to one and the transaction ID stored in the structure is returned to the caller. If the structure is not found, one is created and assigned the next free transaction ID, given a use-count of one and added to the application’s linked list of transactions. The new assigned transaction ID is returned to the caller.

Up-to 64 bytes, including the terminating zero, of the *tran\_name* parameter is saved as the description of the **SpmiCx context** that represents the transaction in the Spmi hierarchy. The key is used as the short name of the context.

### **tran\_detail**

Can be passed in as NULL or some means of specifying a unique instance of the transaction. In the PTX implementation of the ARM API, this parameter is ignored. Consequently, it is not possible to define unique instances of a transaction. If specified as non-NULL, this parameter must be a string not exceeding 128 bytes in length, including the terminating zero.

For the implementation to take this argument in use, another context level would have to be defined between the application context and the transaction context. This was deemed excessive.

### **flags, data, data\_size**

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns an **tran\_id** application identifier. If the subroutine fails, a value less than zero is returned. In compliance with the ARM API specifications, the error return value can be passed to the **arm\_start** (“arm\_start Subroutine” on page 81) subroutine, which will cause **arm\_start** to function as a no-operation.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

`/usr/include/arm.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

`arm_init` (“`arm_init` Subroutine” on page 77) subroutine, `arm_end` (“`arm_end` Subroutine” on page 69) subroutine.

---

## arm\_getid Dual Call Subroutine

### Purpose

The `arm_getid` subroutine is used to register a transaction as belonging to an application and assign a unique identifier to the application/transaction pair. In the PTX implementation of ARM, multiple instances of a transaction within one application can't be defined. The *lower library* implementation of this subroutine may provide support for instances of transactions. A transaction must be registered before any ARM measurements can begin.

### Library

ARM Library (`libarm.a`).

### Syntax

```
#include arm.h

arm_tran_id_t arm_getid(      arm_appl_id_t appl_id,      /* application handle
*/
    arm_ptr_t    *tran_name,    /* transaction name          */
    arm_ptr_t    *tran_detail,  /* transaction additional info */
    arm_flag_t   flags,        /* Reserved = 0              */
    arm_data_t   *data,        /* Reserved = NULL           */
    arm_data_sz_t data_size);  /* Reserved = 0              */
```

### Description

Each transaction needs to be defined by a unique name within an application. Transactions can be defined so they best fit the application environment. For example, if a given environment has thousands of unique transactions, it may be feasible to define groups of similar transactions to prevent data overload. In other situations, you may want to use generated transaction names that reflect what data a transaction carries along with the transaction type. For example, the type of SQL query could be analyzed to group customer query transactions according to complexity, such as *customer\_simple*, *customer*, *customer\_complex*. Whichever method is used to name transactions, in the PTX implementation of the ARM API, measurements are always collected for each unique combination of:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

Before the PTX implementation code is executed, the *lower library* is called. If this call returns a value greater than zero, that return value is passed to the caller as the transaction ID. If the returned value from the *lower library* is zero or negative, the return value is the one generated by the PTX library code.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

Note that the use-count for a transaction structure is either one or zero. This ensures that as long as the application structure is active, so are all transactions for which an **arm\_getid** call was issued after the application was activated by an **arm\_init** (“arm\_init Dual Call Subroutine” on page 79) call. The transaction use-count is reset to zero by the **arm\_end** (“arm\_end Dual Call Subroutine” on page 71) call if this call causes the application use-count to go to zero.

Note that the implementation of **arm\_getid** doesn’t allow unique instances of a transaction to be defined. The **tran\_id** associated with a transaction is stored in the ARM shared memory area and will remain constant throughout the life of the shared memory area. Consequently, subsequent executions of a program that defines one or more transactions under a given application will usually have the same ID returned for the transactions each time. The same is true when different programs define the same transaction within an application: As long as the shared memory area exists, they will all have the same ID returned. This is done to minimize the use of memory for transaction definitions and because it makes no difference from a PTX point of view.

If this is not acceptable from an application point of view, programs can dynamically generate transaction names to pass on the **arm\_getid** subroutine call.

Regardless of the implementation restrictions of the PTX library, the *lower library* may or may not have its own implementation restrictions.

## Parameters

### **appl\_id**

The identifier returned by an earlier call to **arm\_init** (“arm\_init Dual Call Subroutine” on page 79). The identifier is passed to the **arm\_getid** function of the *lower library*. If the *lower library* returns an identifier greater than zero, that identifier is the one that’ll eventually be returned to the caller. After the invocation of the *lower library*, the PTX implementation attempts to translate the *appl\_id* argument to its own identifier by consulting the cross-reference table created by **arm\_init**. If one can be found, it is used for the PTX implementation; if no cross reference is found, the *appl\_id* is used as passed in. The PTX implementation does not require that the **arm\_init** subroutine call was issued by the same program or process now issuing the **arm\_getid** subroutine call. However, the number of issued **arm\_init** subroutine calls for the application name must exceed the number of issued **arm\_end** subroutine calls for this *appl\_id*.

In the PTX implementation, the *appl\_id* (as retrieved from the cross-reference table) is used to look for an application structure. If one is not found or if the use-count of the one found is zero, the PTX implementation is considered to have failed and no action is taken by the PTX library.

### **tran\_name**

A unique transaction name. The name only needs to be unique within the *appl\_id*. The maximum length is 128 characters including the terminating zero. In the PTX implementation, the argument is converted to a key by removing all blanks and truncating the string to 32 characters, including a terminating zero. This key is used to look for a transaction structure (that belongs to the application identified in the first argument) in the library’s private shared memory area. If a transaction structure is found, its use-count is set to one and the transaction ID stored in the structure is saved. If the structure is not found, one is created and assigned the next free transaction ID, given a use-count of one and added to the application’s linked list of transactions. The new assigned transaction ID is saved. If the call to the *lower library* was successful, a cross-reference is created from the *lower library*’s transaction ID to the PTX library’s transaction ID for use by **arm\_start** (“arm\_start Dual Call Subroutine” on page 82).

Up-to 64 bytes, including the terminating zero, of the *tran\_name* parameter is saved as the description of the **SpmiCx context** that represents the transaction in the Spmi hierarchy. The key is used as the short name of the context.

### **tran\_detail**

Can be passed in as NULL or some means of specifying a unique instance of the transaction. In the PTX implementation of the ARM API, this parameter is ignored. Consequently, it is not possible to define unique instances of a transaction. If specified as non-NULL, this parameter must be a string not exceeding 128 bytes in length, including the terminating zero.

For the implementation to take this argument in use, another context level would have to be defined between the application context and the transaction context. This was deemed excessive.

For the *lower library* implementation of this subroutine call, the *tran\_detail* argument may have significance. If so, it's transparent to the PTX implementation.

### **flags, data, data\_size**

In the current API definition, the last three arguments are for future use and they are ignored in the implementation. In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## **Return Values**

If successful, the subroutine returns an **tran\_id** application identifier. If the subroutine fails, a value less than zero is returned. In compliance with the ARM API specifications, the error return value can be passed to the **arm\_start** (“arm\_start Dual Call Subroutine” on page 82) subroutine, which will cause **arm\_start** to function as a no-operation.

If the call to the *lower library* was successful, the **tran\_id** transaction identifier returned is the one assigned by the *lower library*. If the subroutine call to the *lower library* failed but the PTX implementation didn't fail, the **tran\_id** returned is the one assigned by the PTX library. If both implementations fail, a value less than zero is returned. In compliance with the ARM API specification, an error return value can be passed to the **arm\_start** (“arm\_start Dual Call Subroutine” on page 82) subroutine, which will cause **arm\_start** to function as a no-operation.

## **Error Codes**

No error codes are defined by the PTX implementation of the ARM API.

## **Files**

<code>/usr/include/arm.h</code>	Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.
---------------------------------	---

## **Related Information**

**arm\_init** (“arm\_init Dual Call Subroutine” on page 79) subroutine, **arm\_end** (“arm\_end Dual Call Subroutine” on page 71) subroutine.

---

## **arm\_init Subroutine**

### **Purpose**

The **arm\_init** subroutine is used to define an application or a unique instance of an application to the ARM library. In the PTX implementation of ARM, instances of applications can't be defined. An application must be defined before any other ARM subroutine is issued.

## Library

ARM Library (**libarm.a**).

## Syntax

```
#include arm.h

arm_appl_id_t arm_init( arm_ptr_t      *appname,      /* application name
*/
    arm_ptr_t      *appl_user_id, /* Name of the application user */
    arm_flag_t      flags,         /* Reserved = 0 */
    arm_data_t      *data,         /* Reserved = NULL */
    arm_data_sz_t   data_size);   /* Reserved = 0 */
```

## Description

Each application needs to be defined by a unique name. An application can be defined as loosely or as rigidly as required. It may be defined as a single execution of one program, multiple (possibly simultaneous) executions of one program, or multiple executions of multiple programs that together constitute an application. Any one user of ARM may define the application so it best fits the measurement granularity desired. Measurements are always collected for each unique combination of:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

Note that the implementation of **arm\_init** doesn't allow unique instances of an application to be defined. The **appl\_id** associated with an application is stored in the ARM shared memory area and will remain constant throughout the life of the shared memory area. Consequently, subsequent executions of a program that defines one or more applications will usually have the same ID returned for the application each time. The same is true when different programs define the same application: As long as the shared memory area exists, they will all have the same ID returned. This is done to minimize the use of memory for application definitions and because it makes no difference from a PTX point of view.

If this is not acceptable from an application point of view, programs can dynamically generate application names to pass on the **arm\_init** subroutine call.

## Parameters

### **appname**

A unique application name. The maximum length is 128 characters including the terminating zero. The argument is converted to a key by removing all blanks and truncating the string to 32 characters, including a terminating zero. This key is used to look for an application structure in the library's private shared memory area. If a structure is found, its use-count is incremented and the application ID stored in the structure is returned to the caller. If the structure is not found, one is created, assigned the next free application ID and given a use-count of one. The new assigned application ID is returned to the caller.

Up-to 64 bytes, including the terminating zero, of the *appname* parameter is saved as the description of the **SpmiCx context** that represents the application in the Spmi hierarchy. The key is used as the short name of the context.

### **appl\_user\_id**

Can be passed in as NULL or some means of specifying a user ID for the application. This allows the calling program to define unique instances of an application. In the PTX implementation of the

ARM API, this parameter is ignored. Consequently, it is not possible to define unique instances of an application. If specified as non-NULL, this parameter must be a string not exceeding 128 bytes in length, including the terminating zero.

For the implementation to take this argument in use, another context level would have to be defined between the application context and the transaction context. This was deemed excessive.

### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns an **appl\_id** application identifier. If the subroutine fails, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

`/usr/include/arm.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

---

## arm\_init Dual Call Subroutine

### Purpose

The **arm\_init** subroutine is used to define an application or a unique instance of an application to the ARM library. While, in the PTX implementation of ARM, instances of applications can't be defined, the ARM implementation in the *lower library* may permit this. An application must be defined before any other ARM subroutine is issued.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_appl_id_t arm_init( arm_ptr_t    *appname,    /* application name
*/
    arm_ptr_t    *appl_user_id, /* Name of the application user */
    arm_flag_t    flags,        /* Reserved = 0 */
    arm_data_t    *data,        /* Reserved = NULL */
    arm_data_sz_t data_size);   /* Reserved = 0 */
```

### Description

Each application needs to be defined by a unique name. An application can be defined as loosely or as rigidly as required. It may be defined as a single execution of one program, multiple (possibly simultaneous) executions of one program, or multiple executions of multiple programs that together constitute an application. Any one user of ARM may define the application so it best fits the measurement granularity desired. For the PTX implementation, measurements are always collected for each unique combination of:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.

### 3. Unique transaction name.

Before the PTX implementation code is executed, the *lower library* is called. If this call returns a value greater than zero, that return value is passed to the caller as the application ID. If the returned value from the *lower library* is zero or negative, the return value is the one generated by the PTX library code.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

Note that the implementation of **arm\_init** doesn't allow unique instances of an application to be defined. The **appl\_id** associated with an application is stored in the ARM shared memory area and will remain constant throughout the life of the shared memory area. Consequently, subsequent executions of a program that defines one or more applications will usually have the same ID returned for the application each time. The same is true when different programs define the same application: As long as the shared memory area exists, they will all have the same ID returned. This is done to minimize the use of memory for application definitions and because it makes no difference from a PTX point of view.

If this is not acceptable from an application point of view, programs can dynamically generate application names to pass on the **arm\_init** subroutine call.

Regardless of the implementation restrictions of the PTX library, the *lower library* may or may not have its own implementation restrictions.

## Parameters

### **appname**

A unique application name. The maximum length is 128 characters including the terminating zero. The PTX library code converts this value to a key by removing all blanks and truncating the string to 32 characters, including a terminating zero. This key is used to look for an application structure in the library's private shared memory area. If a structure is found, its use-count is incremented and the application ID stored in the structure is saved. If the structure is not found, one is created, assigned the next free application ID and given a use-count of one. The new assigned application ID is saved. If the call to the *lower library* was successful, a cross-reference is created from the *lower library's* application ID to the PTX library's application ID for use by **arm\_getid** ("arm\_getid Dual Call Subroutine" on page 75) and **arm\_end** ("arm\_end Dual Call Subroutine" on page 71).

Up-to 64 bytes, including the terminating zero, of the *appname* parameter is saved as the description of the **SpmiCx context** that represents the application in the Spmi hierarchy. The key is used as the short name of the context.

### **appl\_user\_id**

Can be passed in as NULL or some means of specifying a user ID for the application. This allows the calling program to define unique instances of an application. In the PTX implementation of the ARM API, this parameter is ignored. Consequently, it is not possible to define unique instances of an application. If specified as non-NULL, this parameter must be a string not exceeding 128 bytes in length, including the terminating zero.

For the PTX implementation to take this argument in use, another context level would have to be defined between the application context and the transaction context. This was deemed excessive.

For the *lower library* implementation of this subroutine call, the *appl\_user\_id* argument may have significance. If so, it's transparent to the PTX implementation.

### **flags, data, data\_size**

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.



## Return Values

If the call to the *lower library* was successful, the subroutine returns an **appl\_id** application identifier as returned from the *lower library*. If the subroutine call to the *lower library* fails but the PTX implementation doesn't fail, the **appl\_id** returned is the one assigned by the PTX library. If both implementations fail, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

`/usr/include/arm.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

---

## arm\_start Subroutine

### Purpose

The **arm\_start** subroutine is used to mark the beginning of the execution of a transaction. Measurement of the transaction response time starts at the execution of this subroutine.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_start_handle_t arm_start( arm_tran_id_t tran_id, /* transaction name identifier
*/
    arm_flag_t flags, /* Reserved = 0 */
    arm_data_t *data, /* Reserved = NULL */
    arm_data_sz_t data_size); /* Reserved = 0 */
```

### Description

Each **arm\_start** subroutine call marks the beginning of another instance of a transaction within an application. Multiple instances (simultaneous executions of the transaction) may exist. Control information for the transaction instance is held until the execution of a matching **arm\_stop** (“arm\_stop Subroutine” on page 84) subroutine call, at which time the elapsed time is calculated and used to update transaction measurement metrics for the transaction. Metrics are accumulated for each unique combination of the following three components:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

### Parameters

#### tran\_id

The identifier is returned by an earlier call to **arm\_getid**, “arm\_getid Subroutine” on page 73. The PTX implementation does not require that the **arm\_getid** subroutine call was issued by the same

program or process now issuing the **arm\_start** subroutine call. However, the transaction's application structure must be active, which means that the number of issued **arm\_init** subroutine calls for the application name must exceed the number of issued **arm\_end** subroutine calls for the application's **appl\_id**. If an application was inactivated by issuing a sufficient number of **arm\_end** calls, all transactions defined for that application will have their **use\_count** set to zero. The count remains zero (and the transaction inactive) until a new **arm\_getid** subroutine is issued for the transaction.

The *tran\_id* argument is used to look for a transaction structure. If one is not found or if the use-count of the one found is zero, no action is taken and the function returns -1. If one is found, a transaction instance structure (called a *slot structure*) is allocated, assigned the next free instance ID, and updated with the start time of the transaction instance. The assigned instance ID is returned to the caller.

In compliance with the ARM API specifications, if the *tran\_id* passed is one returned from a previous **arm\_getid** subroutine call that failed, the **arm\_start** subroutine call functions as a no-operation function. It will return a NULL **start\_handle**, which can be passed to subsequent **arm\_update** ("arm\_update Subroutine" on page 88) and **arm\_stop** ("arm\_stop Subroutine" on page 84) subroutine calls with the effect that those calls are no-operation functions.

### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns a **start\_handle**, which uniquely defines this transaction execution instance. If the subroutine fails, a value less than zero is returned. In compliance with the ARM API specifications, the error return value can be passed to the **arm\_update** ("arm\_update Subroutine" on page 88) and **arm\_stop** ("arm\_stop Subroutine" on page 84) subroutines, which will cause those subroutines to operate as no-operation functions.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

<code>/usr/include/arm.h</code>	Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.
---------------------------------	---

## Related Information

**arm\_init** ("arm\_init Subroutine" on page 77) subroutine, **arm\_getid** ("arm\_getid Subroutine" on page 73) subroutine, **arm\_end** ("arm\_end Subroutine" on page 69) subroutine.

---

## arm\_start Dual Call Subroutine

### Purpose

The **arm\_start** subroutine is used to mark the beginning of the execution of a transaction. Measurement of the transaction response time starts at the execution of this subroutine.

### Library

ARM Library (**libarm.a**).

## Syntax

```
#include arm.h

arm_start_handle_t arm_start( arm_tran_id_t tran_id,      /* transaction name identifier
*/
    arm_flag_t    flags,          /* Reserved = 0          */
    arm_data_t    *data,         /* Reserved = NULL      */
    arm_data_sz_t data_size);    /* Reserved = 0          */
```

## Description

Each **arm\_start** subroutine call marks the beginning of another instance of a transaction within an application. Multiple instances (simultaneous executions of the transaction) may exist. Control information for the transaction instance is held until the execution of a matching **arm\_stop** (“arm\_stop Dual Call Subroutine” on page 86) subroutine call, at which time the elapsed time is calculated and used to update transaction measurement metrics for the transaction. Metrics are accumulated for each unique combination of the following three components:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

Before the PTX implementation code is executed, the *lower library* is called. If this call returns a value greater than zero, that return value is passed to the caller as the start handle. If the value returned by the *lower library* is zero or negative, the return value is the one generated by the PTX library code.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

## Parameters

### tran\_id

The identifier is returned by an earlier call to **arm\_getid**, “arm\_getid Dual Call Subroutine” on page 75. The identifier is passed to the **arm\_start** function of the *lower library*. If the *lower library* returns an identifier greater than zero, that identifier is the one that’ll eventually be returned to the caller. After the invocation of the *lower library*, the PTX implementation attempts to translate the *tran\_id* argument to its own identifier from the cross-reference table created by **arm\_getid**. If one can be found, it is used for the PTX implementation; if no cross reference is found, the *tran\_id* is used as passed in. The PTX implementation does not require that the **arm\_getid** subroutine call was issued by the same program or process now issuing the **arm\_start** subroutine call. However, the transaction’s application structure must be active, which means that the number of issued **arm\_init** subroutine calls for the application name must exceed the number of issued **arm\_end** subroutine calls for the application’s *appl\_id*. If an application was inactivated by issuing a sufficient number of **arm\_end** calls, all transactions defined for that application will have their *use\_count* set to zero. The count remains zero (and the transaction inactive) until a new **arm\_getid** subroutine is issued for the transaction.

In the PTX implementation, the *tran\_id* (as retrieved from the cross-reference table) is used to look for a transaction structure. If one is not found or if the use-count of the one found is zero, the PTX implementation is considered to have failed and no action is taken by the PTX library. If one is found, a transaction instance structure (called a *slot structure*) is allocated, assigned the next free instance ID, and updated with the start time of the transaction instance. The assigned instance ID is saved as the **start\_handle**. If the call to the *lower library* was successful, a cross-reference is created from the *lower library*’s *start\_handle* to the PTX library’s *start\_handle* for use by **arm\_update** (“arm\_update Dual Call Subroutine” on page 89) and **arm\_stop** (“arm\_stop Dual Call Subroutine” on page 86).

In compliance with the ARM API specifications, if the *tran\_id* passed is one returned from a previous **arm\_getid** subroutine call that failed, the **arm\_start** subroutine call functions as a no-operation function. It will return a NULL **start\_handle**, which can be passed to subsequent **arm\_update** (“arm\_update Dual Call Subroutine” on page 89) and **arm\_stop** (“arm\_stop Dual Call Subroutine” on page 86) subroutine calls with the effect that those calls are no-operation functions.

### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation. In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns a **start\_handle**, which uniquely defines this transaction execution instance. If the subroutine fails, a value less than zero is returned. In compliance with the ARM API specifications, the error return value can be passed to the **arm\_update** (“arm\_update Dual Call Subroutine” on page 89) and **arm\_stop** (“arm\_stop Dual Call Subroutine” on page 86) subroutines, which will cause those subroutines to operate as no-operation functions.

If the call to the *lower library* was successful, the **start\_handle** instance ID returned is the one assigned by the *lower library*. If the subroutine call to the *lower library* failed but the PTX implementation didn't fail, the **start\_handle** returned is the one assigned by the PTX library. If both implementations fail, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

**/usr/include/arm.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

**arm\_init** (“arm\_init Dual Call Subroutine” on page 79) subroutine, **arm\_getid** (“arm\_getid Dual Call Subroutine” on page 75) subroutine, **arm\_end** (“arm\_end Dual Call Subroutine” on page 71) subroutine.

---

## arm\_stop Subroutine

### Purpose

The **arm\_stop** subroutine is used to mark the end of the execution of a transaction. Measurement of the transaction response time completes at the execution of this subroutine.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_ret_stat_t arm_stop( arm_start_handle_t arm_handle,
                        const arm_status_t comp_status,
                        arm_flag_t flags,
                        arm_data_t * data,
                        arm_data_sz_t data_size);
```

## Description

Each **arm\_stop** subroutine call marks the end of an instance of a transaction within an application. Multiple instances (simultaneous executions of the transaction) may exist. Control information for the transaction instance is held from the execution of the **arm\_start** (“arm\_start Subroutine” on page 81) subroutine call and until the execution of a matching **arm\_stop** subroutine call, at which time the elapsed time is calculated and used to update transaction measurement metrics for the transaction. Metrics are accumulated for each unique combination of the following three components:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

## Parameters

### *arm\_handle*

The identifier is returned by an earlier call to **arm\_start**, “arm\_start Subroutine” on page 81. The *arm\_handle* argument is used to look for a *slot structure* created by the **arm\_start** (“arm\_start Subroutine” on page 81) call, which returned this *arm\_handle*. If one is not found, no action is taken and the function returns -1. If one is found, a *post structure* is allocated and added to the linked list of post structures used to pass data to the **SpmiArmd** daemon. The post structure is updated with the start time from the slot structure, the path to the transaction context, and the stop time of the transaction instance.

In compliance with the ARM API specifications, if the *start\_handle* passed is one returned from a previous **arm\_start** subroutine call that failed, or from an **arm\_start** subroutine operating as a no-operation function, the **arm\_stop** subroutine call executes as a no-operation function. It will return a zero to indicate successful completion.

### **comp\_status**

User supplied transaction completion code. The following codes are defined:

- **ARM\_GOOD** - successful completion. Response time is calculated. The response time is calculated as a fixed point value in milliseconds and saved in the metric **resptime**. In addition, the weighted average response time is calculated as a floating point value using a variable *weight* that defaults to 75%. The average response time is calculated as *weight* percent of the previous value of the average plus (100 - *weight*) percent of the latest response time observation. The value of *weight* can be changed from the **SpmiArmd** daemon’s configuration file **/etc/perf/SpmiArmd.cf**. In addition, the maximum and minimum response time for this transaction is updated, if required. Finally the **count** of successful transaction executions is incremented.
- **ARM\_ABORT** - transaction aborted. The **aborted** counter is incremented. No other updates occur.
- **ARM\_FAILED** - transaction failed. The **failed** counter is incremented. No other updates occur.

### **flags, data, data\_size**

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns zero. If the subroutine fails, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

`/usr/include/arm.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

**arm\_init** (“arm\_init Subroutine” on page 77) subroutine, **arm\_getid** (“arm\_getid Subroutine” on page 73) subroutine, **arm\_start** (“arm\_start Subroutine” on page 81) subroutine, **arm\_end** (“arm\_end Subroutine” on page 69) subroutine.

---

## arm\_stop Dual Call Subroutine

### Purpose

The **arm\_stop** subroutine is used to mark the end of the execution of a transaction. Measurement of the transaction response time completes at the execution of this subroutine.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_ret_stat_t arm_stop( arm_start_handle_t arm_handle, /* unique transaction handle
*/
    const arm_status_t comp_status, /* Good=0, Abort=1, Failed=2 */
    arm_flag_t flags, /* Reserved = 0 */
    arm_data_t *data, /* Reserved = NULL */
    arm_data_sz_t data_size); /* Reserved = 0 */
```

### Description

Each **arm\_stop** subroutine call marks the end of an instance of a transaction within an application. Multiple instances (simultaneous executions of the transaction) may exist. Control information for the transaction instance is held from the execution of the **arm\_start** (“arm\_start Dual Call Subroutine” on page 82) subroutine call and until the execution of a matching **arm\_stop** subroutine call, at which time the elapsed time is calculated and used to update transaction measurement metrics for the transaction. Metrics are accumulated for each unique combination of the following three components:

1. Hostname of the machine where the instrumented application executes.
2. Unique application name.
3. Unique transaction name.

Before the PTX implementation code is executed, the *lower library* is called. If this call returns a value of zero, that return value is passed to the caller. If the value returned by the *lower library* is non-zero, the return value is the one generated by the PTX library code.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product.

## Parameters

### arm\_handle

The identifier is returned by an earlier call to **arm\_start**, “arm\_start Dual Call Subroutine” on page 82. The identifier is passed to the **arm\_stop** function of the *lower library*. If the *lower library* returns a zero return code, that return code is returned to the caller. After the invocation of the *lower library*, the PTX implementation attempts to translate the *arm\_handle* argument to its own identifier from the cross-reference table created by **arm\_start**. If one can be found, it is used for the PTX implementation; if no cross reference is found, the *arm\_handle* is used as passed in. The PTX implementation uses the *start\_handle* argument to look for the *slot structure* created by the **arm\_start** subroutine call. If one is found, a *post structure* is allocated and added to the linked list of post structures used to pass data to the **SpmiArmd** daemon. The post structure is updated with the start time from the slot structure, the path to the transaction context, and the stop time of the transaction instance. If no *slot structure* was found, the PTX implementation is considered to have failed.

In compliance with the ARM API specifications, if the *start\_handle* passed is one returned from a previous **arm\_start** subroutine call that failed, or from an **arm\_start** subroutine operating as a no-operation function, the **arm\_stop** subroutine call executes as a no-operation function. It will return a zero to indicate successful completion.

### comp\_status

User supplied transaction completion code. The following codes are defined:

- **ARM\_GOOD** - successful completion. Response time is calculated. The response time is calculated as a fixed point value in milliseconds and saved in the metric **resptime**. In addition, the weighted average response time (in **respavg**) is calculated as a floating point value using a variable *weight*, that defaults to 75%. The average response time is calculated as *weight* percent of the previous value of the average plus (100 - *weight*) percent of the latest response time observation. The value of *weight* can be changed from the **SpmiArmd** daemon's configuration file **/etc/perf/SpmiArmd.cf**. In addition, the maximum and minimum response time for this transaction is updated, if required. Finally the **count** of successful transaction executions is incremented.
- **ARM\_ABORT** - transaction aborted. The **aborted** counter is incremented. No other updates occur.
- **ARM\_FAILED** - transaction failed. The **failed** counter is incremented. No other updates occur.

### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation. In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns zero. If the subroutine fails, a value less than zero is returned. If the call to the *lower library* was successful, a zero is returned. If the subroutine call to the *lower library* failed but the PTX implementation didn't fail, a zero is returned. If both implementations failed, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

**/usr/include/arm.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

**arm\_init** (“arm\_init Dual Call Subroutine” on page 79) subroutine, **arm\_getid** (“arm\_getid Dual Call Subroutine” on page 75) subroutine, **arm\_start** (“arm\_start Dual Call Subroutine” on page 82) subroutine, **arm\_end** (“arm\_end Dual Call Subroutine” on page 71) subroutine.

---

## arm\_update Subroutine

### Purpose

The **arm\_update** subroutine is used to collect information about a transaction’s progress. It is a no-operation subroutine in the PTX implementation.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_ret_stat_t arm_update( arm_start_handle_t arm_handle, /* unique transaction handle
*/
    arm_flag_t      flags,      /* Reserved = 0          */
    arm_data_t      *data,      /* Reserved = NULL     */
    arm_data_sz_t   data_size); /* Reserved = 0        */
```

### Description

The **arm\_update** subroutine is implemented as a no-operation in the PTX version of the ARM API. It is intended to be used for providing status information for a long-running transaction. Because there’s no feasible way to display such information in current PTX monitors, the subroutine is a NULL function.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product. It is implemented as a NULL subroutine call.

### Parameters

#### start\_handle

The identifier is returned by an earlier call to **arm\_start**, “arm\_start Subroutine” on page 81. The *start\_handle* argument is used to look for the *slot structure* created by the **arm\_start** subroutine call. If one is not found, no action is taken and the function returns -1. Otherwise a zero is returned.

In compliance with the ARM API specifications, if the *start\_handle* passed is one returned from a previous **arm\_start** subroutine call that failed, or from an **arm\_start** subroutine operating as a no-operation function, the **arm\_update** subroutine call executes as a no-operation function. It will return a zero to indicate successful completion.

#### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

### Return Values

If successful, the subroutine returns zero. If the subroutine fails, a value less than zero is returned.

### Error Codes

No error codes are defined by the PTX implementation of the ARM API.



## Files

`/usr/include/arm.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

**arm\_init** (“arm\_init Subroutine” on page 77) subroutine, **arm\_getid** (“arm\_getid Subroutine” on page 73) subroutine, **arm\_start** (“arm\_start Subroutine” on page 81) subroutine, **arm\_stop** (“arm\_stop Subroutine” on page 84) subroutine, **arm\_end** (“arm\_end Subroutine” on page 69) subroutine.

---

## arm\_update Dual Call Subroutine

### Purpose

The **arm\_update** subroutine is used to collect information about a transaction's progress. It is a no-operation subroutine in the PTX implementation but may be fully implemented by the *lower library*.

### Library

ARM Library (**libarm.a**).

### Syntax

```
#include arm.h

arm_ret_stat_t arm_update( arm_start_handle_t arm_handle, /* unique transaction handle
*/
    arm_flag_t      flags,      /* Reserved = 0          */
    arm_data_t      *data,      /* Reserved = NULL     */
    arm_data_sz_t   data_size); /* Reserved = 0          */
```

### Description

The **arm\_update** subroutine is implemented as a no-operation in the PTX version of the ARM API. It is intended to be used for providing status information for a long-running transaction. Because there's no feasible way to display such information in current PTX monitors, the subroutine is a NULL function.

The *lower library* implementation of the **arm\_update** subroutine is always invoked.

This subroutine is part of the implementation of the ARM API in the Performance Toolbox for AIX licensed product. It is implemented as a NULL subroutine call.

### Parameters

#### start\_handle

The identifier is returned by an earlier call to **arm\_start**, “arm\_start Dual Call Subroutine” on page 82. The identifier is passed to the **arm\_update** function of the *lower library*. If the *lower library* returns a zero return code., that return code is returned to the caller. After the invocation of the *lower library*, the PTX implementation attempts to translate the *arm\_handle* argument to its own identifier from the cross-reference table created by **arm\_start**. If one can be found, it is used for the PTX implementation; if no cross reference is found, the *arm\_handle* is used as passed in. The PTX implementation uses the *start\_handle* argument to look for the *slot structure* created by the **arm\_start** subroutine call. If one is found the PTX implementation is considered to have succeeded, otherwise it is considered to have failed.

In compliance with the ARM API specifications, if the *start\_handle* passed is one returned from a previous **arm\_start** subroutine call that failed, or from an **arm\_start** subroutine operating as a no-operation function, the **arm\_update** subroutine call executes as a no-operation function. It will return a zero to indicate successful completion.

### flags, data, data\_size

In the current API definition, the last three arguments are for future use and they are ignored in the implementation. In the current API definition, the last three arguments are for future use and they are ignored in the implementation.

## Return Values

If successful, the subroutine returns zero. If the subroutine fails, a value less than zero is returned. If the call to the *lower library* was successful, a zero is returned. If the subroutine call to the *lower library* failed but the PTX implementation didn't fail, a zero is returned. If both implementations failed, a value less than zero is returned.

## Error Codes

No error codes are defined by the PTX implementation of the ARM API.

## Files

**/usr/include/arm.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the ARM library.

## Related Information

**arm\_init** (“arm\_init Dual Call Subroutine” on page 79) subroutine, **arm\_getid** (“arm\_getid Dual Call Subroutine” on page 75) subroutine, **arm\_start** (“arm\_start Dual Call Subroutine” on page 82) subroutine, **arm\_stop** (“arm\_stop Dual Call Subroutine” on page 86) subroutine, **arm\_end** (“arm\_end Dual Call Subroutine” on page 71) subroutine.

---

## asinh, asinhf, or asinhl Subroutine

### Purpose

Computes the inverse hyperbolic sine.

### Syntax

```
#include <math.h>
```

```
float asinhf (x)  
float x;
```

```
long double asinhl (x)  
long double x;
```

```
double asinh ( x)  
double x;
```

### Description

The **asinhf**, **asinhl**, and **asinh** subroutines compute the inverse hyperbolic sine of the parameter.

An application wishing to check for error situations should set **errno** to zero and call **fetestexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if the **errno** global variable is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **asinhf**, **asinh**, and **asinh** subroutines return the inverse hyperbolic sine of the given argument.

If *x* is NaN, a NaN is returned.

If *x* is 0, or  $\pm\text{Inf}$ , *x* is returned.

If *x* is subnormal, a range error may occur and *x* will be returned.

## Related Information

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## asinf, asinl, or asin Subroutine

### Purpose

Computes the arc sine.

### Syntax

```
#include <math.h>
```

```
float asinf (x)  
float x;
```

```
long double asinl (x)  
long double x;
```

```
double asin (x)  
double x;
```

### Description

The **asinf**, **asinl**, and **asin** subroutines compute the principal value of the arc sine of the *x* parameter. The value of *x* should be in the range [-1,1].

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. On return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **asinf**, **asinl**, and **asin** subroutines return the arc sine of  $x$ , in the range  $[-\pi/2, \pi/2]$  radians.

For finite values of  $x$  not in the range  $[-1, 1]$ , a domain error occurs, and a NaN is returned.

If  $x$  is NaN, a NaN is returned.

If  $x$  is 0,  $x$  is returned.

If  $x$  is  $\pm\text{Inf}$ , a domain error occurs, and a NaN is returned.

If  $x$  is subnormal, a range error may occur and  $x$  is returned.

## Related Information

The “*asinh, asinhf, or asinhl Subroutine*” on page 90.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## assert Macro

### Purpose

Verifies a program assertion.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <assert.h>
```

```
void assert ( Expression )
```

```
int Expression;
```

### Description

The **assert** macro puts error messages into a program. If the specified expression is false, the **assert** macro writes the following message to standard error and stops the program:

```
Assertion failed: Expression, file FileName, line LineNumber
```

In the error message, the *FileName* value is the name of the source file and the *LineNumber* value is the source line number of the **assert** statement.

### Parameters

*Expression* Specifies an expression that can be evaluated as true or false. This expression is evaluated in the same manner as the C language IF statement.

## Related Information

The **abort** (“abort Subroutine” on page 2) subroutine.

The **cpp** command.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## atan2f, atan2l, or atan2 Subroutine

### Purpose

Computes the arc tangent.

### Syntax

```
#include <math.h>
```

```
float atan2f (y, x)  
float y, float x;
```

```
long double atan2l (y, x)  
long double y, long double x;
```

```
double atan2 (y, x)  
double y, x;
```

### Description

The **atan2f**, **atan2l**, and **atan2** subroutines compute the principal value of the arc tangent of  $y/x$ , using the signs of both parameters to determine the quadrant of the return value.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. On return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

<i>y</i>	Specifies the value to compute.
<i>x</i>	Specifies the value to compute.

### Return Values

Upon successful completion, the **atan2f**, **atan2l**, and **atan2** subroutines return the arc tangent of  $y/x$  in the range  $[-\pi, \pi]$  radians.

If  $y$  is 0 and  $x$  is  $< 0$ ,  $\pm\pi$  is returned.

If  $y$  is 0 and  $x$  is  $> 0$ , 0 is returned.

If  $y$  is  $< 0$  and  $x$  is 0,  $-\pi/2$  is returned.

If  $y$  is  $> 0$  and  $x$  is 0,  $\pi/2$  is returned.

If  $x$  is 0, a pole error does not occur.

If either  $x$  or  $y$  is NaN, a NaN is returned.

If the result underflows, a range error may occur and  $y/x$  is returned.

If  $y$  is 0 and  $x$  is  $-0$ ,  $\pm x$  is returned.

If  $y$  is 0 and  $x$  is  $+0$ , 0 is returned.

For finite values of  $\pm y > 0$ , if  $x$  is  $-\text{Inf}$ ,  $\pm x$  is returned.

For finite values of  $\pm y > 0$ , if  $x$  is  $+\text{Inf}$ , 0 is returned.

For finite values of  $x$ , if  $y$  is  $\pm \text{Inf}$ ,  $\pm x/2$  is returned.

If  $y$  is  $\pm \text{Inf}$  and  $x$  is  $-\text{Inf}$ ,  $\pm 3\pi/4$  is returned.

If  $y$  is  $\pm \text{Inf}$  and  $x$  is  $+\text{Inf}$ ,  $\pm \pi/4$  is returned.

If both arguments are 0, a domain error does not occur.

## Related Information

`math.h` in *AIX 5L Version 5.3 Files Reference*.

---

## atan, atanf, or atanl Subroutine

### Purpose

Computes the arc tangent.

### Syntax

```
#include <math.h>
```

```
float atanf (x)  
float x;
```

```
long double atanl (x)  
long double x;
```

```
double atan (x)  
double x;
```

### Description

The `atanf`, `atanl`, and `atan` subroutines compute the principal value of the arc tangent of the  $x$  parameter.

An application wishing to check for error situations should set the `errno` global variable to zero and call `feclearexcept(FE_ALL_EXCEPT)` before calling these functions. On return, if `errno` is nonzero or `fetestexcept(FE_INVALID | FE_DIVBYZERO | FE_OVERFLOW | FE_UNDERFLOW)` is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

## Return Values

Upon successful completion, the **atanf**, **atanl**, and **atan** subroutines return the arc tangent of  $x$  in the range  $[-\pi/2, \pi/2]$  radians.

If  $x$  is NaN, a NaN is returned.

If  $x$  is 0,  $x$  is returned.

If  $x$  is  $\pm\text{Inf}$ ,  $\pm x/2$  is returned.

If  $x$  is subnormal, a range error may occur and  $x$  is returned.

## Related Information

The “atan2f, atan2l, or atan2 Subroutine” on page 93 and “atanh, atanhf, or atanhf Subroutine.”

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## atanh, atanhf, or atanhf Subroutine

### Purpose

Computes the inverse hyperbolic tangent.

### Syntax

```
#include <math.h>
```

```
float atanhf (x)  
float x;
```

```
long double atanhf (x)  
long double x;
```

```
double atanh (x)  
double x;
```

### Description

The **atanhf**, **atanf**, and **atanh** subroutines compute the inverse hyperbolic tangent of the  $x$  parameter.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. On return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

### Return Values

Upon successful completion, the **atanhf**, **atanf**, and **atanh** subroutines return the inverse hyperbolic tangent of the given argument.

If  $x$  is  $\pm 1$ , a pole error occurs, and **atanhf**, **atanf**, and **atanh** return the value of the macro **HUGE\_VALF**, **HUGE\_VALL**, and **HUGE\_VAL** respectively, with the same sign as the correct value of the function.

For finite  $|x| > 1$ , a domain error occurs, and a NaN is returned.

If  $x$  is NaN, a NaN is returned.

If  $x$  is 0,  $x$  is returned.

If  $x$  is  $\pm\text{Inf}$ , a domain error shall occur, and a NaN is returned.

If  $x$  is subnormal, a range error may occur and  $x$  is returned.

## Error Codes

The **atanhf**, **atanhl**, and **atanh** subroutines return **NaNQ** and set **errno** to **EDOM** if the absolute value of  $x$  is greater than 1.

## Related Information

“exp, expf, or expl Subroutine” on page 245

**math.h** in *AIX 5L Version 5.3 Files Reference*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## atof atof Subroutine

### Purpose

Converts an ASCII string to a floating-point or double floating-point number.

### Libraries

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
double atof (NumberPointer)
const char *NumberPointer;
float atof (NumberPointer)
char *NumberPointer;
```

### Description

The **atof** subroutine converts a character string, pointed to by the *NumberPointer* parameter, to a double-precision floating-point number. The **atoff** subroutine converts a character string, pointed to by the *NumberPointer* parameter, to a single-precision floating-point number. The first unrecognized character ends the conversion.

Except for behavior on error, the **atof** subroutine is equivalent to the **strtod** subroutine call, with the *EndPoint* parameter set to (**char\*\***) NULL.

Except for behavior on error, the **atoff** subroutine is equivalent to the **strtof** subroutine call, with the *EndPoint* parameter set to (**char\*\***) NULL.

These subroutines recognize a character string when the characters are in one of two formats: numbers or numeric symbols.

- For a string to be recognized as a number, it should contain the following pieces in the following order:



1. An optional string of white-space characters
  2. An optional sign
  3. A nonempty string of digits optionally containing a radix character
  4. An optional exponent in E-format or e-format followed by an optionally signed integer.
- For a string to be recognized as a numeric symbol, it should contain the following pieces in the following order:
    1. An optional string of white-space characters
    2. An optional sign
    3. One of the strings: **INF**, **infinity**, **NaNQ**, **NaNS**, or **NaN** (case insensitive)

The **atoff** subroutine is not part of the ANSI C Library. These subroutines are at least as accurate as required by the *IEEE Standard for Binary Floating-Point Arithmetic*. The **atof** subroutine accepts at least 17 significant decimal digits. The **atoff** and subroutine accepts at least 9 leading 0's. Leading 0's are not counted as significant digits.

## Parameters

<i>NumberPointer</i>	Specifies a character string to convert.
<i>EndPoint</i>	Specifies a pointer to the character that ended the scan or a null value.

## Return Values

Upon successful completion, the **atof**, and **atoff** subroutines return the converted value. If no conversion could be performed, a value of 0 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

If the conversion cannot be performed, a value of 0 is returned, and the **errno** global variable is set to indicate the error.

If the conversion causes an overflow (that is, the value is outside the range of representable values), **+/- HUGE\_VAL** is returned with the sign indicating the direction of the overflow, and the **errno** global variable is set to **ERANGE**.

If the conversion would cause an underflow, a properly signed value of 0 is returned and the **errno** global variable is set to **ERANGE**.

The **atoff** subroutine has only one rounding error. (If the **atof** subroutine is used to create a double-precision floating-point number and then that double-precision number is converted to a floating-point number, two rounding errors could occur.)

## Related Information

The **scanf** subroutine, **atol**, or **atoi** subroutine, **wstrtol**, **watol**, or **watoi** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## atol or atoll Subroutine

### Purpose

Converts a string to a long integer.

### Syntax

```
#include <stdlib.h>
```

```
long long atoll (nptr)  
const char *nptr;
```

```
long atol (nptr)  
const char *nptr;
```

### Description

The **atoll** and **atol** subroutines (*str*) are equivalent to `strtoll(nptr, (char **)NULL, 10)` and `strtol(nptr, (char **)NULL, 10)`, respectively. If the value cannot be represented, the behavior is undefined.

### Parameters

*nptr*                Points to the string to be converted into a long integer.

### Return Values

The **atoll** and **atol** subroutines return the converted value if the value can be represented.

### Related Information

`strtol`, `strtoul`, `strtoll`, `strtoull`, or `atoi` Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## audit Subroutine

### Purpose

Enables and disables system auditing.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
```

```
int audit ( Command, Argument)  
int Command;  
int Argument;
```

### Description

The **audit** subroutine enables or disables system auditing.

When auditing is enabled, audit records are created for security-relevant events. These records can be collected through the **auditbin** (“auditbin Subroutine” on page 100) subroutine, or through the **/dev/audit** special file interface.

## Parameters

*Command*

Defined in the **sys/audit.h** file, can be one of the following values:

### **AUDIT\_QUERY**

Returns a mask indicating the state of the auditing subsystem. The mask is a logical ORing of the **AUDIT\_ON**, **AUDIT\_OFF**, and **AUDIT\_PANIC** flags. The *Argument* parameter is ignored.

### **AUDIT\_ON**

Enables auditing. If auditing is already enabled, only the failure-mode behavior changes. The *Argument* parameter specifies recovery behavior in the event of failure and may be either 0 or the value **AUDIT\_PANIC**.

**Note:** If **AUDIT\_PANIC** is specified, bin-mode auditing must be enabled before the **audit** subroutine call.

### **AUDIT\_OFF**

Disables the auditing system if auditing is enabled. If the auditing system is disabled, the **audit** subroutine does nothing. The *Argument* parameter is ignored.

### **AUDIT\_RESET**

Disables the auditing system (as does **AUDIT\_OFF**) and resets the auditing system. If auditing is already disabled, only the system configuration is reset. Resetting the audit configuration involves clearing the audit events and audited objects table, and terminating bin and stream auditing. The *Argument* parameter is ignored.

### **AUDIT\_EVENT\_THRESHOLD**

Audit event records will be buffered until a total of *Argument* records have been saved, at which time the audit event records will be flushed to disk. An *Argument* value of zero disables this functionality. This parameter only applies to AIX 4.1.4 and later.

### **AUDIT\_BYTE\_THRESHOLD**

Audit event data will be buffered until a total of *Argument* bytes of data have been saved, at which time the audit event data will be flushed to disk. An *Argument* value of zero disables this functionality. This parameter only applies to AIX 4.1.4 and later.

*Argument*

Specifies the behavior when a bin write fails (for **AUDIT\_ON**) or specifies the size of the audit event buffer (for **AUDIT\_EVENT\_THRESHOLD** and **AUDIT\_BYTE\_THRESHOLD**). For all other commands, the value of **Argument** is ignored. The valid values are:

### **AUDIT\_PANIC**

The operating system halts abruptly if an audit record cannot be written to a bin.

**Note:** If **AUDIT\_PANIC** is specified, bin-mode auditing must be enabled before the **audit** subroutine call.

### **BufferSize**

The number of bytes or audit event records which will be buffered. This parameter is valid only with the command **AUDIT\_BYTE\_THRESHOLD** and **AUDIT\_EVENT\_THRESHOLD**. A value of zero will disable either byte (for **AUDIT\_BYTE\_THRESHOLD**) or event (for **AUDIT\_EVENT\_THRESHOLD**) buffering.

## Return Values

For a *Command* value of **AUDIT\_QUERY**, the **audit** subroutine returns, upon successful completion, a mask indicating the state of the auditing subsystem. The mask is a logical ORing of the **AUDIT\_ON**, **AUDIT\_OFF**, **AUDIT\_PANIC**, and **AUDIT\_NO\_PANIC** flags. For any other *Command* value, the **audit** subroutine returns 0 on successful completion.

If the **audit** subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **audit** subroutine fails if either of the following is true:

<b>EINVAL</b>	The <i>Command</i> parameter is not one of <b>AUDIT_ON</b> , <b>AUDIT_OFF</b> , <b>AUDIT_RESET</b> , or <b>AUDIT_QUERY</b> .
<b>EINVAL</b>	The <i>Command</i> parameter is <b>AUDIT_ON</b> and the <i>Argument</i> parameter specifies values other than <b>AUDIT_PANIC</b> .
<b>EPERM</b>	The calling process does not have root user authority.

## Files

**dev/audit** Specifies the audit pseudo-device from which the audit records are read.

## Related Information

The **auditbin** (“auditbin Subroutine”) subroutine, **auditevents** (“auditevents Subroutine” on page 102) subroutine, **auditlog** (“auditlog Subroutine” on page 104) subroutine, **auditobj** (“auditobj Subroutine” on page 105) subroutine, **auditproc** (“auditproc Subroutine” on page 109) subroutine.

The **audit** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditbin Subroutine

### Purpose

Defines files to contain audit records.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
```

```
int auditbin (Command, Current, Next, Threshold)
int Command;
int Current;
int Next;
int Threshold;
```

### Description

The **auditbin** subroutine establishes an audit bin file into which the kernel writes audit records. Optionally, this subroutine can be used to establish an overflow bin into which records are written when the current bin reaches the size specified by the *Threshold* parameter.

## Parameters

<i>Command</i>	<p>If nonzero, this parameter is a logical ORing of the following values, which are defined in the <b>sys/audit.h</b> file:</p> <p><b>AUDIT_EXCL</b> Requests exclusive rights to the audit bin files. If the file specified by the <i>Current</i> parameter is not the kernel's current bin file, the <b>auditbin</b> subroutine fails immediately with the <b>errno</b> variable set to <b>EBUSY</b>.</p> <p><b>AUDIT_WAIT</b> The <b>auditbin</b> subroutine should not return until:</p> <p><b>bin full</b> The kernel writes the number of bytes specified by the <i>Threshold</i> parameter to the file descriptor specified by the <i>Current</i> parameter. Upon successful completion, the <b>auditbin</b> subroutine returns a 0. The kernel writes subsequent audit records to the file descriptor specified by the <i>Next</i> parameter.</p> <p><b>bin failure</b> An attempt to write an audit record to the file specified by the <i>Current</i> parameter fails. If this occurs, the <b>auditbin</b> subroutine fails with the <b>errno</b> variable set to the return code from the <b>auditwrite</b> subroutine.</p> <p><b>bin contention</b> Another process has already issued a successful call to the <b>auditbin</b> subroutine. If this occurs, the <b>auditbin</b> subroutine fails with the <b>errno</b> variable set to <b>EBUSY</b>.</p> <p><b>system shutdown</b> The auditing system was shut down. If this occurs, the <b>auditbin</b> subroutine fails with the <b>errno</b> variable set to <b>EINTR</b>.</p>
<i>Current</i>	A file descriptor for a file to which the kernel should immediately write audit records.
<i>Next</i>	Specifies the file descriptor that will be used as the current audit bin if the value of the <i>Threshold</i> parameter is exceeded or if a write to the current bin fails. If this value is -1, no switch occurs.
<i>Threshold</i>	Specifies the maximum size of the current bin. If 0, the auditing subsystem will not switch bins. If it is nonzero, the kernel begins writing records to the file specified by the <i>Next</i> parameter, if writing a record to the file specified by the <i>Cur</i> parameter would cause the size of this file to exceed the number of bytes specified by the <i>Threshold</i> parameter. If no next bin is defined and <b>AUDIT_PANIC</b> was specified when the auditing subsystem was enabled, the system is shut down. If the size of the <i>Threshold</i> parameter is too small to contain a bin header and a bin tail, the <b>auditbin</b> subroutine fails and the <b>errno</b> variable is set to <b>EINVAL</b> .

## Return Values

If the **auditbin** subroutine is successful, a value of 0 returns.

If the **auditbin** subroutine fails, a value of -1 returns and the **errno** global variable is set to indicate the error. If this occurs, the result of the call does not indicate whether any records were written to the bin.

## Error Codes

The **auditbin** subroutine fails if any of the following is true:

<b>EBADF</b>	The <i>Current</i> parameter is not a file descriptor for a regular file open for writing, or the <i>Next</i> parameter is neither -1 nor a file descriptor for a regular file open for writing.
<b>EBUSY</b>	The <i>Command</i> parameter specifies <b>AUDIT_EXCL</b> and the kernel is not writing audit records to the file specified by the <i>Current</i> parameter.
<b>EBUSY</b>	The <i>Command</i> parameter specifies <b>AUDIT_WAIT</b> and another process has already registered a bin.
<b>EINTR</b>	The auditing subsystem is shut down.

<b>EINVAL</b>	The <i>Command</i> parameter specifies a nonzero value other than <b>AUDIT_EXCL</b> or <b>AUDIT_WAIT</b> .
<b>EINVAL</b>	The <i>Threshold</i> parameter value is less than the size of a bin header and trailer.
<b>EPERM</b>	The caller does not have root user authority.

## Related Information

The **audit** (“audit Subroutine” on page 98) subroutine, **auditevents** (“auditevents Subroutine”) subroutine, **auditlog** (“auditlog Subroutine” on page 104) subroutine, **auditobj** (“auditobj Subroutine” on page 105) subroutine, **auditproc** (“auditproc Subroutine” on page 109) subroutine.

The **audit** command.

The **audit** file format.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditevents Subroutine

### Purpose

Gets or sets the status of system event auditing.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
```

```
int auditevents ( Command, Classes, NClasses)
int Command;
struct audit_class *Classes;
int NClasses;
```

### Description

The **auditevents** subroutine queries or sets the audit class definitions that control event auditing. Each audit class is a set of one or more audit events.

System auditing need not be enabled before calling the **auditevents** subroutine. The **audit** (“audit Subroutine” on page 98) subroutine can be directed with the **AUDIT\_RESET** command to clear all event lists.

## Parameters

<i>Command</i>	Specifies whether the event lists are to be queried or set. The values, defined in the <b>sys/audit.h</b> file, for the <i>Command</i> parameter are:  <b>AUDIT_SET</b> Sets the lists of audited events after first clearing all previous definitions.  <b>AUDIT_GET</b> Queries the lists of audited events.  <b>AUDIT_LOCK</b> Queries the lists of audited events. This value also blocks any other process attempting to set or lock the list of audit events. The lock is released when the process holding the lock dies or calls the <b>auditevents</b> subroutine with the <i>Command</i> parameter set to <b>AUDIT_SET</b> .
<i>Classes</i>	Specifies the array of <b>a_event</b> structures for the <b>AUDIT_SET</b> operation, or after an <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> operation. The <b>audit_class</b> structure is defined in the <b>sys/audit.h</b> file and contains the following members:  <b>ae_name</b> A pointer to the name of the audit class.  <b>ae_list</b> A pointer to a list of null-terminated audit event names for this audit class. The list is ended by a null name (a leading null byte or two consecutive null bytes). <b>Note:</b> Event and class names are limited to 15 significant characters.  <b>ae_len</b> The length of the event list in the <b>ae_list</b> member. This length includes the terminating null bytes. On an <b>AUDIT_SET</b> operation, the caller must set this member to indicate the actual length of the list (in bytes) pointed to by <b>ae_list</b> . On an <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> operation, the <b>auditevents</b> subroutine sets this member to indicate the actual size of the list.
<i>NClasses</i>	Serves a dual purpose. For <b>AUDIT_SET</b> , the <i>NClasses</i> parameter specifies the number of elements in the events array. For <b>AUDIT_GET</b> and <b>AUDIT_LOCK</b> , the <i>NClasses</i> parameter specifies the size of the buffer pointed to by the <i>Classes</i> parameter.

**Attention:** Only 32 audit classes are supported. One class is implicitly defined by the system to include all audit events (ALL). The administrator of your system should not attempt to define more than 31 audit classes.

## Security

The calling process must have root user authority in order to use the **auditevents** subroutine.

## Return Codes

If the **auditevents** subroutine completes successfully, the number of audit classes is returned if the *Command* parameter is **AUDIT\_GET** or **AUDIT\_LOCK**. A value of 0 is returned if the *Command* parameter is **AUDIT\_SET**. If this call fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **auditevents** subroutine fails if one or more of the following are true:

<b>EPERM</b>	The calling process does not have root user authority.
<b>EINVAL</b>	The value of <i>Command</i> is not <b>AUDIT_SET</b> , <b>AUDIT_GET</b> , or <b>AUDIT_LOCK</b> .
<b>EINVAL</b>	The <i>Command</i> parameter is <b>AUDIT_SET</b> , and the value of the <i>NClasses</i> parameter is greater than or equal to 32.
<b>EINVAL</b>	A class name or event name is longer than 15 significant characters.

<b>ENOSPC</b>	The value of <i>Command</i> is <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> and the size of the buffer specified by the <i>NClasses</i> parameter is not large enough to hold the list of event structures and names. If this occurs, the first word of the buffer is set to the required buffer size.
<b>EFAULT</b>	The <i>Classes</i> parameter points outside of the process' address space.
<b>EFAULT</b>	The <i>ae_list</i> member of one or more <b>audit_class</b> structures passed for an <b>AUDIT_SET</b> operation points outside of the process' address space.
<b>EFAULT</b>	The <i>Command</i> value is <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> and the size of the <i>Classes</i> buffer is not large enough to hold an integer.
<b>EBUSY</b>	Another process has already called the <b>auditevents</b> subroutine with <b>AUDIT_LOCK</b> .
<b>ENOMEM</b>	Memory allocation failed.

## Related Information

The **audit** (“audit Subroutine” on page 98) subroutine, **auditbin** (“auditbin Subroutine” on page 100) subroutine, **auditlog** (“auditlog Subroutine”) subroutine, **auditobj** (“auditobj Subroutine” on page 105) subroutine, **auditproc** (“auditproc Subroutine” on page 109) subroutine, **auditread** (“auditread, auditread\_r Subroutines” on page 111) subroutine, **auditwrite** (“auditwrite Subroutine” on page 112) subroutine.

The **audit** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditlog Subroutine

### Purpose

Appends an audit record to the audit trail file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
```

```
int auditlog ( Event, Result, Buffer, BufferSize)
char *Event;
int Result;
char *Buffer;
int BufferSize;
```

### Description

The **auditlog** subroutine generates an audit record. The kernel audit-logging component appends a record for the specified *Event* if system auditing is enabled, process auditing is not suspended, and the *Event* parameter is in one or more of the audit classes for the current process.

The audit logger generates the audit record by adding the *Event* and *Result* parameters to the audit header and including the resulting information in the *Buffer* parameter as the audit tail.

### Parameters

*Event* The name of the audit event to be generated. This parameter should be the name of an audit event. Audit event names are truncated to 15 characters plus null.



<i>Result</i>	Describes the result of this event. Valid values are defined in the <b>sys/audit.h</b> file and include the following: <b>AUDIT_OK</b> The event was successful. <b>AUDIT_FAIL</b> The event failed. <b>AUDIT_FAIL_ACCESS</b> The event failed because of any access control denial. <b>AUDIT_FAIL_DAC</b> The event failed because of a discretionary access control denial. <b>AUDIT_FAIL_PRIV</b> The event failed because of a privilege control denial. <b>AUDIT_FAIL_AUTH</b> The event failed because of an authentication denial.
<i>Buffer</i>	Other nonzero values of the <i>Result</i> parameter are converted into the <b>AUDIT_FAIL</b> value. Points to a buffer containing the tail of the audit record. The format of the information in this buffer depends on the event name.
<i>BufferSize</i>	Specifies the size of the <i>Buffer</i> parameter, including the terminating null.

## Return Values

Upon successful completion, the **auditlog** subroutine returns a value of 0. If **auditlog** fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

The **auditlog** subroutine does not return any indication of failure to write the record where this is due to inappropriate tailoring of auditing subsystem configuration files or user-written code. Accidental omissions and typographical errors in the configuration are potential causes of such a failure.

## Error Codes

The **auditlog** subroutine fails if any of the following are true:

<b>EFAULT</b>	The <i>Event</i> or <i>Buffer</i> parameter points outside of the process' address space.
<b>EINVAL</b>	The auditing system is either interrupted or not initialized.
<b>EINVAL</b>	The length of the audit record is greater than 32 kilobytes.
<b>EPERM</b>	The process does not have root user authority.
<b>ENOMEM</b>	Memory allocation failed.

## Related Information

The **audit** (“audit Subroutine” on page 98) subroutine, **auditbin** (“auditbin Subroutine” on page 100) subroutine, **auditevents** (“auditevents Subroutine” on page 102) subroutine, **auditobj** (“auditobj Subroutine”) subroutine, **auditproc** (“auditproc Subroutine” on page 109) subroutine, **auditwrite** (“auditwrite Subroutine” on page 112) subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditobj Subroutine

### Purpose

Gets or sets the auditing mode of a system data object.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/audit.h>
```

```
int auditobj ( Command, Obj_Events, ObjSize)
int Command;
struct o_event *Obj_Events;
int ObjSize;
```

## Description

The **auditobj** subroutine queries or sets the audit events to be generated by accessing selected objects. For each object in the file system name space, it is possible to specify the event generated for each access mode. Using the **auditobj** subroutine, an administrator can define new audit events in the system that correspond to accesses to specified objects. These events are treated the same as system-defined events.

System auditing need not be enabled to set or query the object audit events. The **audit** subroutine can be directed with the **AUDIT\_RESET** command to clear the definitions of object audit events.

## Parameters

*Command*

Specifies whether the object audit event lists are to be read or written. The valid values, defined in the **sys/audit.h** file, for the *Command* parameter are:

**AUDIT\_SET**

Sets the list of object audit events, after first clearing all previous definitions.

**AUDIT\_GET**

Queries the list of object audit events.

**AUDIT\_LOCK**

Queries the list of object audit events and also blocks any other process attempting to set or lock the list of audit events. The lock is released when the process holding the lock dies or calls the **auditobj** subroutine with the *Command* parameter set to **AUDIT\_SET**.

## Obj\_Events

Specifies the array of **o\_event** structures for the **AUDIT\_SET** operation or for after the **AUDIT\_GET** or **AUDIT\_LOCK** operation. The **o\_event** structure is defined in the **sys/audit.h** file and contains the following members:

**o\_type** Specifies the type of the object, in terms of naming space. Currently, only one object-naming space is supported:

### **AUDIT\_FILE**

Denotes the file system naming space.

**o\_name** Specifies the name of the object.

### **o\_event**

Specifies any array of event names to be generated when the object is accessed. Note that event names are currently limited to 16 bytes, including the trailing null. The index of an event name in this array corresponds to an access mode. Valid indexes are defined in the **audit.h** file and include the following:

- **AUDIT\_READ**
- **AUDIT\_WRITE**
- **AUDIT\_EXEC**

**Note:** The C++ compiler will generate a warning indicating that **o\_event** is defined both as a structure and a field within that structure. Although the **o\_event** field can be used within C++, the warning can be bypassed by defining **O\_EVENT\_RENAME**. This will replace the **o\_event** field with **o\_event\_array**. **o\_event** is the default field.

## ObjSize

For an **AUDIT\_SET** operation, the *ObjSize* parameter specifies the number of object audit event definitions in the array pointed to by the *Obj\_Events* parameter. For an **AUDIT\_GET** or **AUDIT\_LOCK** operation, the *ObjSize* parameter specifies the size of the buffer pointed to by the *Obj\_Events* parameter.

## Return Values

If the **auditobj** subroutine completes successfully, the number of object audit event definitions is returned if the *Command* parameter is **AUDIT\_GET** or **AUDIT\_LOCK**. A value of 0 is returned if the *Command* parameter is **AUDIT\_SET**. If this call fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **auditobj** subroutine fails if any of the following are true:

<b>EFAULT</b>	The <i>Obj_Events</i> parameter points outside the address space of the process.
<b>EFAULT</b>	The <i>Command</i> parameter is <b>AUDIT_SET</b> , and one or more of the <i>o_name</i> members points outside the address space of the process.
<b>EFAULT</b>	The <i>Command</i> parameter is <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> , and the buffer size of the <i>Obj_Events</i> parameter is not large enough to hold the integer.
<b>EINVAL</b>	The value of the <i>Command</i> parameter is not <b>AUDIT_SET</b> , <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> .
<b>EINVAL</b>	The <i>Command</i> parameter is <b>AUDIT_SET</b> , and the value of one or more of the <i>o_type</i> members is not <b>AUDIT_FILE</b> .
<b>EINVAL</b>	An event name was longer than 15 significant characters.
<b>ENOENT</b>	The <i>Command</i> parameter is <b>AUDIT_SET</b> , and the parent directory of one of the file-system objects does not exist.
<b>ENOSPC</b>	The value of the <i>Command</i> parameter is <b>AUDIT_GET</b> or <b>AUDIT_LOCK</b> , and the size of the buffer as specified by the <i>ObjSize</i> parameter is not large enough to hold the list of event structures and names. If this occurs, the first word of the buffer is set to the required buffer size.
<b>ENOMEM</b>	Memory allocation failed.
<b>EBUSY</b>	Another process has called the <b>auditobj</b> subroutine with <b>AUDIT_LOCK</b> .

**EPERM** The caller does not have root user authority.

## Related Information

The **audit** (“audit Subroutine” on page 98) subroutine, **auditbin** (“auditbin Subroutine” on page 100) subroutine, **auditevents** (“auditevents Subroutine” on page 102) subroutine, **auditlog** (“auditlog Subroutine” on page 104) subroutine, **auditproc** (“auditproc Subroutine” on page 109) subroutine.

The **audit** command.

The **audit.h** file.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditpack Subroutine

### Purpose

Compresses and uncompresses audit bins.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
#include <stdio.h>
```

```
char *auditpack ( Expand, Buffer)
int Expand;
char *Buffer;
```

### Description

The **auditpack** subroutine can be used to compress or uncompress bins of audit records.

### Parameters

*Expand* Specifies the operation. Valid values, as defined in the **sys/audit.h** header file, are one of the following:

**AUDIT\_PACK**

Performs standard compression on the audit bin.

**AUDIT\_UNPACK**

Unpacks the compressed audit bin.

*Buffer* Specifies the buffer containing the bin to be compressed or uncompressed. This buffer must contain a standard bin as described in the **audit.h** file.

### Return Values

If the **auditpack** subroutine is successful, a pointer to a buffer containing the processed audit bin is returned. If unsuccessful, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **auditpack** subroutine fails if one or more of the following values is true:

<b>EINVAL</b>	The <i>Expand</i> parameter is not one of the valid values ( <b>AUDIT_PACK</b> or <b>AUDIT_UNPACK</b> ).
<b>EINVAL</b>	The <i>Expand</i> parameter is <b>AUDIT_UNPACK</b> and the packed data in <i>Buffer</i> does not unpack to its original size.
<b>EINVAL</b>	The <i>Expand</i> parameter is <b>AUDIT_PACK</b> and the bin in the <i>Buffer</i> parameter is already compressed, or the <i>Expand</i> parameter is <b>AUDIT_UNPACK</b> and the bin in the <i>Buffer</i> parameter is already unpacked.
<b>ENOSPC</b>	The <b>auditpack</b> subroutine is unable to allocate space for a new buffer.

## Related Information

The **auditread** (“auditread, auditread\_r Subroutines” on page 111) subroutine.

The **auditcat** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditproc Subroutine

### Purpose

Gets or sets the audit state of a process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
```

```
int auditproc (ProcessID, Command, Argument, Length)
int ProcessID;
int Command;
char * Argument;
int Length;
```

### Description

The **auditproc** subroutine queries or sets the auditing state of a process. There are two parts to the auditing state of a process:

- The list of classes to be audited for this process. Classes are defined by the **auditevents** (“auditevents Subroutine” on page 102) subroutine. Each class includes a set of audit events. When a process causes an audit event, that event may be logged in the audit trail if it is included in one or more of the audit classes of the process.
- The audit status of the process. Auditing for a process may be suspended or resumed. Functions that generate an audit record can first check to see whether auditing is suspended. If process auditing is suspended, no audit events are logged for a process. For more information, see the **auditlog** (“auditlog Subroutine” on page 104) subroutine.

## Parameters

<i>ProcessID</i>	The process ID of the process to be affected. If <i>ProcessID</i> is 0, the <b>auditproc</b> subroutine affects the current process.
<i>Command</i>	The action to be taken. Defined in the <b>audit.h</b> file, valid values include: <b>AUDIT_KLIST_EVENTS</b> Sets the list of audit classes to be audited for the process and also sets the user's default audit classes definition within the kernel. The <i>Argument</i> parameter is a pointer to a list of null-terminated audit class names. The <i>Length</i> parameter is the length of this list, including null bytes. <b>AUDIT_QEVENTS</b> Returns the list of audit classes defined for the current process if <i>ProcessID</i> is 0. Otherwise, it returns the list of audit classes defined for the specified process ID. The <i>Argument</i> parameter is a pointer to a character buffer. The <i>Length</i> parameter specifies the size of this buffer. On return, this buffer contains a list of null-terminated audit class names. A null name terminates the list. <b>AUDIT_EVENTS</b> Sets the list of audit classes to be audited for the process. The <i>Argument</i> parameter is a pointer to a list of null-terminated audit class names. The <i>Length</i> parameter is the length of this list, including null bytes. <b>AUDIT_QSTATUS</b> Returns the audit status of the current process. You can only check the status of the current process. If the <i>ProcessID</i> parameter is nonzero, a -1 is returned and the <b>errno</b> global variable is set to <b>EINVAL</b> . The <i>Length</i> and <i>Argument</i> parameters are ignored. A return value of <b>AUDIT_SUSPEND</b> indicates that auditing is suspended. A return value of <b>AUDIT_RESUME</b> indicates normal auditing for this process. <b>AUDIT_STATUS</b> Sets the audit status of the current process. The <i>Length</i> parameter is ignored, and the <i>ProcessID</i> parameter must be zero. If <i>Argument</i> is <b>AUDIT_SUSPEND</b> , the audit status is set to suspend event auditing for this process. If the <i>Argument</i> parameter is <b>AUDIT_RESUME</b> , the audit status is set to resume event auditing for this process.
<i>Argument</i>	A character pointer for the audit class buffer for an <b>AUDIT_EVENT</b> or <b>AUDIT_QEVENTS</b> value of the <i>Command</i> parameter or an integer defining the audit status to be set for an <b>AUDIT_STATUS</b> operation.
<i>Length</i>	Size of the audit class character buffer.

## Return Values

The **auditproc** subroutine returns the following values upon successful completion:

- The previous audit status (**AUDIT\_SUSPEND** or **AUDIT\_RESUME**), if the call queried or set the audit status (the *Command* parameter specified **AUDIT\_QSTATUS** or **AUDIT\_STATUS**)
- A value of 0 if the call queried or set audit events (the *Command* parameter specified **AUDIT\_QEVENTS** or **AUDIT\_EVENTS**)

## Error Codes

If the **auditproc** subroutine fails if one or more of the following are true:

<b>EINVAL</b>	An invalid value was specified for the <i>Command</i> parameter.
<b>EINVAL</b>	The <i>Command</i> parameter is set to the <b>AUDIT_QSTATUS</b> or <b>AUDIT_STATUS</b> value and the <b>pid</b> value is nonzero.
<b>EINVAI</b>	The <i>Command</i> parameter is set to the <b>AUDIT_STATUS</b> value and the <i>Argument</i> parameter is not set to <b>AUDIT_SUSPEND</b> or <b>AUDIT_RESUME</b> .
<b>ENOSPC</b>	The <i>Command</i> parameter is <b>AUDIT_QEVENTS</b> , and the buffer size is insufficient. In this case, the first word of the <i>Argument</i> parameter is set to the required size.

<b>EFAULT</b>	The <i>Command</i> parameter is <b>AUDIT_QEVENTS</b> or <b>AUDIT_EVENTS</b> and the <i>Argument</i> parameter points to a location outside of the process' allocated address space.
<b>ENOMEM</b>	Memory allocation failed.
<b>EPERM</b>	The caller does not have root user authority.

## Related Information

The **audit** (“audit Subroutine” on page 98) subroutine, **auditbin** (“auditbin Subroutine” on page 100) subroutine, **auditevents** (“auditevents Subroutine” on page 102) subroutine, **auditlog** (“auditlog Subroutine” on page 104) subroutine, **auditobj** (“auditobj Subroutine” on page 105) subroutine, **auditwrite** (“auditwrite Subroutine” on page 112) subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditread, auditread\_r Subroutines

### Purpose

Reads an audit record.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
#include <stdio.h>
char *auditread ( FilePointer, AuditRecord)
FILE *FilePointer;
struct aud_rec *AuditRecord;

char *auditread_r ( FilePointer, AuditRecord, RecordSize, StreamInfo)
FILE *FilePointer;
struct aud_rec *AuditRecord;
size_t RecordSize;
void **StreamInfo;
```

### Description

The **auditread** subroutine reads the next audit record from the specified file descriptor. Bins on this input stream are unpacked and uncompressed if necessary.

The **auditread** subroutine can not be used on more than one *FilePointer* as the results can be unpredictable. Use the **auditread\_r** subroutine instead.

The **auditread\_r** subroutine reads the next audit from the specified file descriptor. This subroutine is thread safe and can be used to handle multiple open audit files simultaneously by multiple threads of execution.

The **auditread\_r** subroutine is able to read multiple versions of audit records. The version information contained in an audit record is used to determine the correct size and format of the record. When an input record header is larger than *AuditRecord*, an error is returned. In order to provide for binary compatibility with previous versions, if *RecordSize* is the same size as the original (**struct aud\_rec**), the input record is converted to the original format and returned to the caller.

## Parameters

<i>FilePointer</i>	Specifies the file descriptor from which to read.
<i>AuditRecord</i>	Specifies the buffer to contain the header. The first short in this buffer must contain a valid number for the header.
<i>RecordSize</i>	The size of the buffer referenced by <i>AuditRecord</i> .
<i>StreamInfo</i>	A pointer to an opaque datatype used to hold information related to the current value of <i>FilePointer</i> . For each new value of <i>FilePointer</i> , a new <i>StreamInfo</i> pointer must be used. <i>StreamInfo</i> must be initialized to NULL by the user and is initialized by <b>auditread_r</b> when first used. When <i>FilePointer</i> has been closed, the value of <i>StreamInfo</i> can be passed to the <b>free</b> subroutine to be deallocated.

## Return Values

If the **auditread** subroutine completes successfully, a pointer to a buffer containing the tail of the audit record is returned. The length of this buffer is returned in the `ah_length` field of the header file. If this subroutine is unsuccessful, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **auditread** subroutine fails if one or more of the following is true:

<b>EBADF</b>	The <i>FilePointer</i> value is not valid.
<b>ENOSPC</b>	The <b>auditread</b> subroutine is unable to allocate space for the tail buffer.

Other error codes are returned by the **read** subroutine.

## Related Information

The **auditpack** (“auditpack Subroutine” on page 108) subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## auditwrite Subroutine

### Purpose

Writes an audit record.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <sys/audit.h>
#include <stdio.h>
```

```
int auditwrite (Event, Result, Buffer1, Length1, Buffer2, Length2, ...)
char * Event;
int Result;
char * Buffer1, *Buffer2 ...;
```

```
int Length1, Length2 ...;
```



## Description

The **auditwrite** subroutine builds the tail of an audit record and then writes it with the **auditlog** subroutine. The tail is built by gathering the specified buffers. The last buffer pointer must be a null.

If the **auditwrite** subroutine is to be called from a program invoked from the **inittab** file, the **setpcred** subroutine should be called first to establish the process' credentials.

## Parameters

<i>Event</i>	Specifies the name of the event to be logged.
<i>Result</i>	Specifies the audit status of the event. Valid values are defined in the <b>sys/audit.h</b> file and are listed in the <b>auditlog</b> subroutine.
<i>Buffer1, Buffer2</i>	Specifies the character buffers containing audit tail information. Note that numerical values must be passed by reference. The correct size can be computed with the <b>sizeof</b> C function.
<i>Length1, Length2</i>	Specifies the lengths of the corresponding buffers.

## Return Values

If the **auditwrite** subroutine completes successfully, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **auditwrite** subroutine fails if the following is true:

**ENOSPC**            The **auditwrite** subroutine is unable to allocate space for the tail buffer.

Other error codes are returned by the **auditlog** subroutine.

## Related Information

The **auditlog** ("auditlog Subroutine" on page 104) subroutine, **setpcred** subroutine.

The **inittab** file.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## authenticate Subroutine

### Purpose

Verifies a user's name and password.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int authenticate (UserName, Response, Reenter, Message)
char *UserName;
char *Response;
int *Reenter;
char **Message;
```

## Description

The **authenticate** subroutine maintains requirements users must satisfy to be authenticated to the system. It is a callable interface that prompts for the user's name and password. The user must supply a character string at the prompt issued by the *Message* parameter. The *Response* parameter returns the user's response to the **authenticate** subroutine. The calling program makes no assumptions about the number of prompt messages the user must satisfy for authentication.

The *Reenter* parameter indicates when a user has satisfied all prompt messages. The parameter remains nonzero until a user has passed all prompts. After the returned value of *Reenter* is 0, the return code signals whether authentication has succeeded or failed. When progressing through prompts for a user, the value of *Reenter* must be maintained by the caller between invocations of **authenticate**.

The **authenticate** subroutine ascertains the authentication domains the user can attempt. The subroutine reads the **SYSTEM** line from the user's stanza in the */etc/security/user* file. Each token that appears in the **SYSTEM** line corresponds to a method that can be dynamically loaded and processed. Likewise, the system can provide multiple or alternate authentication paths.

The **authenticate** routine maintains internal state information concerning the next prompt message presented to the user. If the calling program supplies a different user name before all prompts are complete for the user, the internal state information is reset and prompt messages begin again. The calling program maintains the value of the *Reenter* parameter while processing prompts for a given user.

If the user has no defined password, or the **SYSTEM** grammar explicitly specifies no authentication required, the user is not required to respond to any prompt messages. Otherwise, the user is always initially prompted to supply a password.

The **authenticate** subroutine can be called initially with the cleartext password in the *Response* parameter. If the user supplies a password during the initial invocation but does not have a password, authentication fails. If the user wants the **authenticate** subroutine to supply a prompt message, the *Response* parameter is a null pointer on initial invocation.

The **authenticate** subroutine sets the **AUTHSTATE** environment variable used by name resolution subroutines, such as the **getpwnam** subroutine. This environment variable indicates the registry to which to user authenticated. Values for the **AUTHSTATE** environment variable include **DCE**, **compat**, and token names that appear in a **SYSTEM** grammar. A null value can exist if the **cron** daemon or other utilities that do not require authentication is called.

## Parameters

<i>UserName</i>	Points to the user's name that is to be authenticated.
<i>Response</i>	Specifies a character string containing the user's response to an authentication prompt.
<i>Reenter</i>	Points to a Boolean value that signals whether the <b>authenticate</b> subroutine has completed processing. If the <i>Reenter</i> parameter is a nonzero value, the <b>authenticate</b> subroutine expects the user to satisfy the prompt message provided by the <i>Message</i> parameter. If the <i>Reenter</i> parameter is 0, the <b>authenticate</b> subroutine has completed processing.
<i>Message</i>	Points to a pointer that the <b>authenticate</b> subroutine allocates memory for and fills in. This string is suitable for printing and issues prompt messages (if the <i>Reenter</i> parameter is a nonzero value). It also issues informational messages such as why the user failed authentication (if the <i>Reenter</i> parameter is 0). The calling application is responsible for freeing this memory.

## Return Values

Upon successful completion, the **authenticate** subroutine returns a value of 0. If this subroutine fails, it returns a value of 1.

## Error Codes

The **authenticate** subroutine is unsuccessful if one of the following values is true:

<b>ENOENT</b>	Indicates that the user is unknown to the system.
<b>ESAD</b>	Indicates that authentication is denied.
<b>EINVAL</b>	Indicates that the parameters are not valid.
<b>ENOMEM</b>	Indicates that memory allocation (malloc) failed.

**Note:** The DCE mechanism requires credentials on successful authentication that apply only to the authenticate process and its children.

## Related Information

The **ckuserID** (“ckuserID Subroutine” on page 166) subroutine.

---

## authenticate Subroutine

### Purpose

Verifies a user’s name and password.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int authenticate (UserName, Response, Reenter, Message, State)
char *UserName;
char *Response;
int *Reenter;
char **Message;
void **State;
```

### Description

The **authenticate** subroutine maintains requirements that users must satisfy to be authenticated to the system. It is a callable interface that prompts for the user’s name and password. The user must supply a character string at the prompt issued by the *Message* parameter. The *Response* parameter returns the user’s response to the **authenticate** subroutine. The calling program makes no assumptions about the number of prompt messages the user must satisfy for authentication. The **authenticate** subroutine maintains information about the results of each part of the authentication process in the *State* parameter. This parameter can be shared with the **chpassx**, **loginrestrictionsx** and **passwdexpiredx** subroutines. The proper sequence of library routines for authenticating a user in order to create a new session is:

1. Call the **loginrestrictionsx** subroutine to determine which administrative domains allow the user to log in.
2. Call the **authenticate** subroutine to perform authentication using those administrative domains that grant login access.
3. Call the **passwdexpiredx** subroutine to determine if any of the passwords used during the authentication process have expired and must be changed in order for the user to be granted access.
4. If the **passwdexpiredx** subroutine indicated that one or more passwords have expired and must be changed by the user, call the **chpassx** subroutine to update all of the passwords that were used for the authentication process.

The *Reenter* parameter remains a nonzero value until the user satisfies all prompt messages or answers incorrectly. When the *Reenter* parameter is 0, the return code signals whether authentication passed or failed. The value of the *Reenter* parameter must be 0 on the initial call. A nonzero value for the *Reenter* parameter must be passed to the **authenticate**x subroutine on subsequent calls. A new authentication can be begun by calling the **authenticate**x subroutine with a 0 value for the *Reenter* parameter or by using a different value for *UserName*.

The *State* parameter contains information about the authentication process. The *State* parameter from an earlier call to **loginrestrictions**x can be used to control how authentication is performed. Administrative domains that do not permit the user to log in cause those administrative domains to be ignored during authentication even if the user has the correct authentication information.

The **authenticate**x subroutine ascertains the authentication domains the user can attempt. The subroutine uses the **SYSTEM** attribute for the user. Each token that is displayed in the **SYSTEM** line corresponds to a method that can be dynamically loaded and processed. Likewise, the system can provide multiple or alternate authentication paths.

The **authenticate**x subroutine maintains internal state information concerning the next prompt message presented to the user. If the calling program supplies a different user name before all prompts are complete for the user, the internal state information is reset and prompt messages begin again. The **authenticate**x subroutine requires that the *State* parameter be initialized to reference a null value when changing user names or that the *State* parameter from an earlier call to **loginrestrictions**x for the new user be provided.

If the user has no defined password, or the **SYSTEM** grammar explicitly specifies no authentication required, the user is not required to respond to any prompt messages. Otherwise, the user is always initially prompted to supply a password.

The **authenticate**x subroutine can be called initially with the *cleartext* password in the *Response* parameter. If the user supplies a password during the initial invocation but does not have a password, authentication fails. If the user wants the **authenticate**x subroutine to supply a prompt message, the *Response* parameter is a null pointer on initial invocation.

The **authenticate**x subroutine sets the **AUTHSTATE** environment variable used by name resolution subroutines, such as the **getpwnam** subroutine. This environment variable indicates the first registry to which the user authenticated. Values for the **AUTHSTATE** environment variable include **DCE**, **compat**, and token names that appear in a **SYSTEM** grammar. A null value can exist if the **cron** daemon or another utility that does not require authentication is called.

## Parameters

<i>Message</i>	Points to a pointer that the <b>authenticate</b> x subroutine allocates memory for and fills in. This string is suitable for printing and issues prompt messages (if the <i>Reenter</i> parameter is a nonzero value). It also issues informational messages, such as why the user failed authentication (if the <i>Reenter</i> parameter is 0). The calling application is responsible for freeing this memory.
<i>Reenter</i>	Points to an integer value that signals whether the <b>authenticate</b> x subroutine has completed processing. If the integer referenced by the <i>Reenter</i> parameter is a nonzero value, the <b>authenticate</b> x subroutine expects the user to satisfy the prompt message provided by the <i>Message</i> parameter. If the integer referenced by the <i>Reenter</i> parameter is 0, the <b>authenticate</b> x subroutine has completed processing. The initial value of the integer referenced by <i>Reenter</i> must be 0 when the <b>authenticate</b> x function is initially invoked and must not be modified by the calling application until the <b>authentication</b> x subroutine has completed processing.
<i>Response</i>	Specifies a character string containing the user's response to an authentication prompt.

<i>State</i>	Points to a pointer that the <b>authenticatex</b> subroutine allocates memory for and fills in. The <i>State</i> parameter can also be the result of an earlier call to the <b>loginrestrictionsx</b> subroutine. This parameter contains information about the results of the authentication process for each term in the user's <b>SYSTEM</b> attribute. The calling application is responsible for freeing this memory when it is no longer needed for a subsequent call to the <b>passwdexpiredx</b> or <b>chpassx</b> subroutines.
<i>UserName</i>	Points to the user's name that is to be authenticated.

## Return Values

Upon successful completion, the **authenticatex** subroutine returns a value of 0. If this subroutine fails, it returns a value of 1.

## Error Codes

The **authenticatex** subroutine is unsuccessful if one of the following values is true:

<b>EINVAL</b>	The parameters are not valid.
<b>ENOENT</b>	The user is unknown to the system.
<b>ENOMEM</b>	Memory allocation (malloc) failed.
<b>ESAD</b>	Authentication is denied.

**Note:** Additional information about the behavior of a loadable authentication module can be found in the documentation for that module.

## Related Information

The “authenticate Subroutine” on page 113, “chpassx Subroutine” on page 156, “loginrestrictionsx Subroutine” on page 754, “passwdexpiredx Subroutine” on page 974.

---

## basename Subroutine

### Purpose

Return the last element of a path name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <libgen.h>
```

```
char *basename (char *path)
```

### Description

Given a pointer to a character string that contains a path name, the **basename** subroutine deletes trailing “/” characters from *path*, and then returns a pointer to the last component of *path*. The “/” character is defined as trailing if it is not the first character in the string.

If *path* is a null pointer or points to an empty string, a pointer to a static constant “.” is returned.

### Return Values

The **basename** function returns a pointer to the last component of *path*.

The **basename** function returns a pointer to a static constant "." if *path* is a null pointer or points to an empty string.

The **basename** function may modify the string pointed to by *path* and may return a pointer to static storage that may then be overwritten by a subsequent call to the **basename** subroutine.

## Examples

Input string	Output string
"/usr/lib"	"lib"
"/usr/"	"usr"
"/"	"/"

## Related Information

The **dirname** ("dirname Subroutine" on page 213) subroutine.

---

## bcopy, bcmp, bzero or ffs Subroutine

### Purpose

Performs bit and byte string operations.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <strings.h>
```

```
void bcopy (Source, Destination, Length)
```

```
const void *Source,
```

```
char *Destination;
```

```
size_t Length;
```

```
int bcmp (String1, String2, Length)
```

```
const void *String1, *String2;
```

```
size_t Length;
```

```
void bzero (String,Length)
```

```
char *String;
```

```
int Length;
```

```
int ffs (Index)
```

```
int Index;
```

### Description

**Note:** The **bcopy** subroutine takes parameters backwards from the **strcpy** subroutine.

The **bcopy**, **bcmp**, and **bzero** subroutines operate on variable length strings of bytes. They do not check for null bytes as do the **string** routines.

The **bcopy** subroutine copies the value of the *Length* parameter in bytes from the string in the *Source* parameter to the string in the *Destination* parameter.

The **bcmp** subroutine compares the byte string in the *String1* parameter against the byte string of the *String2* parameter, returning a zero value if the two strings are identical and a nonzero value otherwise. Both strings are assumed to be *Length* bytes long.

The **bzero** subroutine zeroes out the string in the *String* parameter for the value of the *Length* parameter in bytes.

The **ffs** subroutine finds the first bit set in the *Index* parameter passed to it and returns the index of that bit. Bits are numbered starting at 1. A return value of 0 indicates that the value passed is 0.

## Related Information

The **memcmp**, **memccpy**, **memchr**, **memcpy**, **memmove**, **memset** (“memccpy, memchr, memcmp, memcpy, memset or memmove Subroutine” on page 807) subroutines, **strcat**, **strncat**, **strxfrm**, **strcpy**, **strncpy**, or **strdup** subroutine, **strcmp**, **strncmp**, **strcasecmp**, **strncasecmp**, or **strcoll** subroutine, **strlen**, **strchr**, **strchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, or **strtok** subroutine, **swab** subroutine.

List of String Manipulation Subroutines and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## bessel: j0, j1, jn, y0, y1, or yn Subroutine

### Purpose

Computes Bessel functions.

### Libraries

IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
```

```
double j0 (x)
```

```
double x;
```

```
double j1 (x)
```

```
double x;
```

```
double jn (n, x)
```

```
int n;
```

```
double x;
```

```
double y0 (x)
```

```
double x;
```

```
double y1 (x)
```

```
double x;
```

```
double yn (n, x)
```

```
int n;
```

```
double x;
```

### Description

Bessel functions are used to compute wave variables, primarily in the field of communications.

The **j0** subroutine and **j1** subroutine return Bessel functions of *x* of the first kind, of orders 0 and 1, respectively. The **jn** subroutine returns the Bessel function of *x* of the first kind of order *n*.

The **y0** subroutine and **y1** subroutine return the Bessel functions of  $x$  of the second kind, of orders 0 and 1, respectively. The **yn** subroutine returns the Bessel function of  $x$  of the second kind of order  $n$ . The value of  $x$  must be positive.

**Note:** Compile any routine that uses subroutines from the **libm.a** library with the **-lm** flag. To compile the **j0.c** file, for example:

```
cc j0.c -lm
```

## Parameters

$x$  Specifies some double-precision floating-point value.  
 $n$  Specifies some integer value.

## Return Values

When using **libm.a** (**-lm**), if  $x$  is negative, **y0**, **y1**, and **yn** return the value NaNQ. If  $x$  is 0, **y0**, **y1**, and **yn** return the value **-HUGE\_VAL**.

When using **libmsaa.a** (**-lmsaa**), values too large in magnitude cause the functions **j0**, **j1**, **y0**, and **y1** to return 0 and to set the **errno** global variable to ERANGE. In addition, a message indicating TLOSS error is printed on the standard error output.

Nonpositive values cause **y0**, **y1**, and **yn** to return the value **-HUGE** and to set the **errno** global variable to EDOM. In addition, a message indicating argument DOMAIN error is printed on the standard error output.

These error-handling procedures may be changed with the **matherr** subroutine when using **libmsaa.a** (**-lmsaa**).

## Related Information

The **matherr** (“matherr Subroutine” on page 789) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## bindprocessor Subroutine

### Purpose

Binds kernel threads to a processor.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/processor.h>
```

```
int bindprocessor ( What, Who, Where )  
int What;  
int Who;  
cpu_t Where;
```



## Description

The **bindprocessor** subroutine binds a single kernel thread, or all kernel threads in a process, to a processor, forcing the bound threads to be scheduled to run on that processor. It is important to understand that a process itself is not bound, but rather its kernel threads are bound. Once kernel threads are bound, they are always scheduled to run on the chosen processor, unless they are later unbound. When a new thread is created, it has the same bind properties as its creator. This applies to the initial thread in the new process created by the **fork** subroutine: the new thread inherits the bind properties of the thread which called **fork**. When the **exec** subroutine is called, thread properties are left unchanged.

The **bindprocessor** subroutine will fail if the target process has a *Resource Attachment*.

Programs that use processor bindings should become Dynamic Logical Partitioning (DLPAR) aware. Refer to Dynamic Logical Partitioning in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for more information.

## Parameters

<i>What</i>	Specifies whether a process or a thread is being bound to a processor. The <i>What</i> parameter can take one of the following values: <b>BINDPROCESS</b> A process is being bound to a processor. <b>BINDTHREAD</b> A thread is being bound to a processor.
<i>Who</i>	Indicates a process or thread identifier, as appropriate for the <i>What</i> parameter, specifying the process or thread which is to be bound to a processor.
<i>Where</i>	If the <i>Where</i> parameter is a bind CPU identifier, it specifies the processor to which the process or thread is to be bound. A value of <b>PROCESSOR_CLASS_ANY</b> unbinds the specified process or thread, which will then be able to run on any processor.  The <b>sysconf</b> subroutine can be used to retrieve information about the number of online processors in the system.

## Return Values

On successful completion, the **bindprocessor** subroutine returns 0. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **bindprocessor** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	The <i>What</i> parameter is invalid, or the <i>Where</i> parameter indicates an invalid processor number or a processor class which is not currently available.
<b>ESRCH</b>	The specified process or thread does not exist.
<b>EPERM</b>	The caller does not have root user authority, and the <i>Who</i> parameter specifies either a process, or a thread belonging to a process, having a real or effective user ID different from that of the calling process. The target process has a <i>Resource Attachment</i> .

## Related Information

The **bindprocessor** command.

The **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exact Subroutine” on page 236) subroutine, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **sysconf** subroutine, **thread\_self** subroutine.

## brk or sbrk Subroutine

### Purpose

Changes data segment space allocation.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int brk ( EndDataSegment )  
char *EndDataSegment;
```

```
void *sbrk ( Increment )  
intptr_t Increment;
```

### Description

The **brk** and **sbrk** subroutines dynamically change the amount of space allocated for the data segment of the calling process. (For information about segments, see the **exec** subroutine. For information about the maximum amount of space that can be allocated, see the **ulimit** and **getrlimit** subroutines.)

The change is made by resetting the break value of the process, which determines the maximum space that can be allocated. The break value is the address of the first location beyond the current end of the data region. The amount of available space increases as the break value increases. The available space is initialized to a value of 0 at the time it is used. The break value can be automatically rounded up to a size appropriate for the memory management architecture.

The **brk** subroutine sets the break value to the value of the *EndDataSegment* parameter and changes the amount of available space accordingly.

The **sbrk** subroutine adds to the break value the number of bytes contained in the *Increment* parameter and changes the amount of available space accordingly. The *Increment* parameter can be a negative number, in which case the amount of available space is decreased.

### Parameters

<i>EndDataSegment</i>	Specifies the effective address of the maximum available data.
<i>Increment</i>	Specifies any integer.

### Return Values

Upon successful completion, the **brk** subroutine returns a value of 0, and the **sbrk** subroutine returns the old break value. If either subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **brk** subroutine and the **sbrk** subroutine are unsuccessful and the allocated space remains unchanged if one or more of the following are true:

<b>ENOMEM</b>	The requested change allocates more space than is allowed by a system-imposed maximum. (For information on the system-imposed maximum on memory space, see the <b>ulimit</b> system call.)
<b>ENOMEM</b>	The requested change sets the break value to a value greater than or equal to the start address of any attached shared-memory segment. (For information on shared memory operations, see the <b>shmat</b> subroutine.)

## Related Information

The **exec** (“**exec**: **execl**, **execle**, **execlp**, **execv**, **execve**, **execvp**, or **exec** Subroutine” on page 236) subroutines, **getrlimit** (“**getrlimit**, **getrlimit64**, **setrlimit**, **setrlimit64**, or **vlimit** Subroutine” on page 424) subroutine, **shmat** subroutine, **shmdt** subroutine, **ulimit** subroutine.

The **\_end** (“**\_end**, **\_etext**, or **\_edata** Identifier” on page 224), **\_etext** (“**\_end**, **\_etext**, or **\_edata** Identifier” on page 224), or **\_edata** (“**\_end**, **\_etext**, or **\_edata** Identifier” on page 224) identifier.

Subroutine Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## bsearch Subroutine

### Purpose

Performs a binary search.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
void *bsearch ( Key, Base, NumberOfElements, Size, ComparisonPointer)
const void *Key;
const void *Base;
size_t NumberOfElements;
size_t Size;
int (*ComparisonPointer) (const void *, const void *);
```

### Description

The **bsearch** subroutine is a binary search routine.

The **bsearch** subroutine searches an array of *NumberOfElements* objects, the initial member of which is pointed to by the *Base* parameter, for a member that matches the object pointed to by the *Key* parameter. The size of each member in the array is specified by the *Size* parameter.

The array must already be sorted in increasing order according to the provided comparison function *ComparisonPointer* parameter.

### Parameters

<i>Key</i>	Points to the object to be sought in the array.
<i>Base</i>	Points to the element at the base of the table.
<i>NumberOfElements</i>	Specifies the number of elements in the array.

<i>ComparisonPointer</i>	Points to the comparison function, which is called with two arguments that point to the <i>Key</i> parameter object and to an array member, in that order.
<i>Size</i>	Specifies the size of each member in the array.

## Return Values

If the *Key* parameter value is found in the table, the **bsearch** subroutine returns a pointer to the element found.

If the *Key* parameter value is not found in the table, the **bsearch** subroutine returns the null value. If two members compare as equal, the matching member is unspecified.

For the *ComparisonPointer* parameter, the comparison function compares its parameters and returns a value as follows:

- If the first parameter is less than the second parameter, the *ComparisonPointer* parameter returns a value less than 0.
- If the first parameter is equal to the second parameter, the *ComparisonPointer* parameter returns a value of 0.
- If the first parameter is greater than the second parameter, the *ComparisonPointer* parameter returns a value greater than 0.

The comparison function need not compare every byte, so arbitrary data can be contained in the elements in addition to the values being compared.

The *Key* and *Base* parameters should be of type pointer-to-element and cast to type pointer-to-character. Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

## Related Information

The **hsearch** (“hsearch, hcreate, or hdestroy Subroutine” on page 527) subroutine, **lsearch** (“lsearch or lfind Subroutine” on page 764) subroutine, **qsort** subroutine.

Knuth, Donald E.; *The Art of Computer Programming*, Volume 3. Reading, Massachusetts, Addison-Wesley, 1981.

Searching and Sorting Example Program and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

## btowc Subroutine

### Purpose

Single-byte to wide-character conversion.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <stdio.h>
#include <wchar.h>
wint_t btowc (intc);
```

## Description

The *btowc* function determines whether *c* constitutes a valid (one-byte) character in the initial shift state.

The behavior of this function is affected by the LC\_CTYPE category of the current locale.

## Return Values

The *btowc* function returns WEOF if *c* has the value EOF or if (unsigned char) *c* does not constitute a valid (one-byte) character in the initial shift state. Otherwise, it returns the wide-character representation of that character.

## Related Information

The *wctob* subroutine.

---

## buildproclist Subroutine

### Purpose

Retrieves a list of process transaction records based on the criteria specified.

### Library

The **libaacct.a** library.

### Syntax

```
#define <sys/aacct.h>
int buildproclist(crit, crit_list, n_crit, p_list, sublist)
int crit;
union proc_crit *crit_list;
int n_crit;
struct aaacct_tran_rec *p_list;
struct aaacct_tran_rec **sublist;
```

### Description

The **buildproclist** subroutine retrieves a subset of process transaction records from the master process transaction records that are given as input based on the selection criteria provided. This selection criteria can be one of the following values defined in **sys/aacct.h**:

- **CRIT\_UID**
- **CRIT\_GID**
- **CRIT\_PROJ**
- **CRIT\_CMD**

For example, if the criteria is specified as **CRIT\_UID**, the list of process transaction records for specific user IDs will be retrieved. The list of user IDs are passed through the *crit\_list* argument of type **union proc\_crit**. Based on the specified criteria, the caller has to pass an array of user IDs, group IDs, project IDs or command names in this union.

Usually, the master list of transaction records is obtained by a prior call to the **getproclist** subroutine.

### Parameters

<i>crit</i>	Integer value representing the selection criteria for the process records.
<i>crit_list</i>	Pointer to <b>union proc_crit</b> where the data for the selection criteria is passed.
<i>n_crit</i>	Number of elements to be considered for the selection, such as the number of user IDs.
<i>p_list</i>	Master list of process transaction records.

*sublist* Pointer to the linked list of **aacct\_tran\_rec** structures, which hold the retrieved process transaction records.

## Security

No restrictions. Any user can call this function.

## Return Values

0 The call to the subroutine was successful.  
-1 The call to the subroutine failed.

## Error Codes

**EINVAL** The passed pointer is NULL.  
**ENOMEM** Insufficient memory.  
**EPERM** Permission denied. Unable to read the data file.

## Related Information

The “buildproclist Subroutine” on page 125, “buildtranlist or freetranlist Subroutine,” “getfilehdr Subroutine” on page 367.

The acctprt command.

Understanding the Advanced Accounting Subsystem.

---

## buildtranlist or freetranlist Subroutine

### Purpose

Read the advanced accounting records from the advanced accounting data file.

### Library

The **libaacct.a** library.

### Syntax

```
#define <sys/aacct.h>
buildtranlist(filename, trid[], ntrids, begin_time, end_time, tran_list)
char *filename;
unsigned int trid[];
unsigned int ntrids;
long long begin_time;
long long end_time;
struct aacct_tran_rec **tran_list;
freetranlist(tran_list)
struct aacct_tran_rec *tran_list;
```

### Description

The **buildtranlist** subroutine retrieves the transaction records of the specified transaction type from the accounting data file. The required transaction IDs are passed as arguments, and these IDs are defined in **sys/aacct.h**. The list of transaction records are returned to the calling program through the *tran\_list* pointer argument.

This API can be called multiple times with different accounting data file names to generate a consolidated list of transaction records from multiple data files. It appends the new file data to the end of the linked list pointed to by the *tran\_list* argument. In addition, it internally sorts the transaction records based on the time of transaction so users can get a time-sorted list of transaction records from this routine. This subroutine can also be used to retrieve the intended transaction records for a particular interval of time by specifying the begin and end times of this interval as arguments.

The **freetrانlist** subroutine frees the memory allocated to these transaction records. It can be used to deallocate memory that has been allocated to the transaction record lists created by routines such as **buildtranlist**, **getproclist**, **getlparlist**, and **getarmlist**.

## Parameters

<i>begin_time</i>	Specifies the start timestamp for collecting records in a particular intervals. The input is in seconds since EPOCH. Specifying -1 retrieves all the records.
<i>end_time</i>	Specifies the end timestamp for collecting records in a particular intervals. The input is in seconds since EPOCH. Specifying -1 retrieves all the records.
<i>filename</i>	Name of the advanced accounting data file.
<i>ntrids</i>	Count of transaction IDs passed in the array <i>trid</i> .
<i>tran_list</i>	Pointer to the linked list of <b>aacct_tran_rec</b> structures that are to be returned to the caller or freed.
<i>trid</i>	An array of transaction record type identifiers.

## Security

No restrictions. Any user can call this function.

## Return Values

0	The call to the subroutine was successful.
-1	The call to the subroutine failed.

## Error Codes

<b>EINVAL</b>	The passed pointer is NULL.
<b>ENOENT</b>	Specified data file does not exist.
<b>ENOMEM</b>	Insufficient memory.
<b>EPERM</b>	Permission denied. Unable to read the data file.

## Related Information

The “buildproclist Subroutine” on page 125, “getproclist, getlparlist, or getarmlist Subroutine” on page 414.

Understanding the Advanced Accounting Subsystem.

---

## **\_check\_lock Subroutine**

### **Purpose**

Conditionally updates a single word variable atomically.

### **Library**

Standard C library (**libc.a**)

## Syntax

```
#include <sys/atomic_op.h>
```

```
boolean_t _check_lock ( word_addr, old_val, new_val)
atomic_p word_addr;
int old_val;
int new_val;
```

## Parameters

<i>word_addr</i>	Specifies the address of the single word variable.
<i>old_val</i>	Specifies the old value to be checked against the value of the single word variable.
<i>new_val</i>	Specifies the new value to be conditionally assigned to the single word variable.

## Description

The **\_check\_lock** subroutine performs an atomic (uninterruptible) sequence of operations. The **compare\_and\_swap** subroutine is similar, but does not issue synchronization instructions and therefore is inappropriate for updating lock words.

**Note:** The word variable must be aligned on a full word boundary.

## Return Values

<b>FALSE</b>	Indicates that the single word variable was equal to the old value and has been set to the new value.
<b>TRUE</b>	Indicates that the single word variable was not equal to the old value and has been left unchanged.

## Related Information

The **\_clear\_lock** (“\_clear\_lock Subroutine”) subroutine.

---

## **\_clear\_lock** Subroutine

### Purpose

Stores a value in a single word variable atomically.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/atomic_op.h>
```

```
void _clear_lock ( word_addr, value)
atomic_p word_addr;
int value
```

### Parameters

<i>word_addr</i>	Specifies the address of the single word variable.
<i>value</i>	Specifies the value to store in the single word variable.



## Description

The `_clear_lock` subroutine performs an atomic (uninterruptible) sequence of operations.

This subroutine has no return values.

**Note:** The word variable must be aligned on a full word boundary.

## Related Information

The `_check_lock` (“\_check\_lock Subroutine” on page 127) subroutine.

---

## cabs, cabsf, or cabsl Subroutine

### Purpose

Returns a complex absolute value.

### Syntax

```
#include <complex.h>
```

```
double cabs (z)
double complex z;
```

```
float cabsf (z)
float complex z;
```

```
long double cabsl (z)
long double complex z;
```

### Description

The `cabs`, `cabsf`, or `cabsl` subroutines compute the complex absolute value (also called norm, modulus, or magnitude) of the `z` parameter.

### Parameters

`z` Specifies the value to be computed.

### Return Values

Returns the complex absolute value.

---

## cacos, cacosf, or cacosl Subroutine

### Purpose

Computes the complex arc cosine.

### Syntax

```
#include <complex.h>
```

```
double complex cacos (z)
double complex z;
```

```
float complex cacosf (z)
```

```
float complex z;
```

```
long double complex cacosl (z)  
long double complex z;
```

## Description

The **cacos**, **cacosf**, or **cacosl** subroutine computes the complex arc cosine of  $z$ , with branch cuts outside the interval  $[-1, +1]$  along the real axis.

## Parameters

$z$  Specifies the value to be computed.

## Return Values

The **cacos**, **cacosf**, or **cacosl** subroutine returns the complex arc cosine value, in the range of a strip mathematically unbounded along the imaginary axis and in the interval  $[0, \pi]$  along the real axis.

---

## cacosh, cacoshf, or cacoshl Subroutines

### Purpose

Computes the complex arc hyperbolic cosine.

### Syntax

```
#include <complex.h>
```

```
double complex cacosh (z)  
double complex z;
```

```
float complex cacoshf (z)  
float complex z;
```

```
long double complex cacoshl (z)  
long double complex z;
```

### Description

The **cacosh**, **cacoshf**, or **cacoshl** subroutine computes the complex arc hyperbolic cosine of the  $z$  parameter, with a branch cut at values less than 1 along the real axis.

### Parameters

$z$  Specifies the value to be computed.

### Return Values

The **cacosh**, **cacoshf**, or **cacoshl** subroutine returns the complex arc hyperbolic cosine value, in the range of a half-strip of non-negative values along the real axis and in the interval  $[-i\pi, +i\pi]$  along the imaginary axis.

### Related Information

The “ccosh, ccoshf, or ccoshl Subroutine” on page 139.

---

## carg, cargf, or cargl Subroutine

### Purpose

Returns the complex argument value.

### Syntax

```
#include <complex.h>
```

```
double carg (z)
double complex z;
```

```
float cargf (z)
float complex z;
```

```
long double cargl (z)
long double complex z;
```

### Description

The **carg**, **cargf**, or **cargl** subroutine computes the argument (also called phase angle) of the *z* parameter, with a branch cut along the negative real axis.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **carg**, **cargf**, or **cargl** subroutine returns the value of the argument in the interval  $[-\pi, +\pi]$ .

### Related Information

The “**cimag**, **cimagf**, or **cimagl** Subroutine” on page 163, “**conj**, **conjf**, or **conjl** Subroutine” on page 182, and “**cproj**, **cprojf**, or **cprojl** Subroutine” on page 190.

---

## casin, casinf, or casinl Subroutine

### Purpose

Computes the complex arc sine.

### Syntax

```
#include <complex.h>
```

```
double complex casin (z)
double complex z;
```

```
float complex casinf (z)
float complex z;
```

```
long double complex casinl (z)
long double complex z;
```

### Description

The **casin**, **casinf**, or **casinl** subroutine computes the complex arc sine of the *z* parameter, with branch cuts outside the interval  $[-1, +1]$  along the real axis.

## Parameters

*z* Specifies the value to be computed.

## Return Values

The **casin**, **casinf**, or **casinl** subroutine returns the complex arc sine value, in the range of a strip mathematically unbounded along the imaginary axis and in the interval  $[-\pi/2, +\pi/2]$  along the real axis.

## Related Information

The “**csin**, **csinf**, or **csinl** Subroutine” on page 193.

---

## casinh, casinfh, or casinlh Subroutine

### Purpose

Computes the complex arc hyperbolic sine.

### Syntax

```
#include <complex.h>
```

```
double complex casinh (z)
double complex z;
```

```
float complex casinhf (z)
float complex z;
```

```
long double complex casinhl (z)
long double complex z;
```

### Description

The **casinh**, **casinfh**, and **casinlh** subroutines compute the complex arc hyperbolic sine of the *z* parameter, with branch cuts outside the interval  $[-i, +i]$  along the imaginary axis.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **casinh**, **casinfh**, and **casinlh** subroutines return the complex arc hyperbolic sine value, in the range of a strip mathematically unbounded along the real axis and in the interval  $[-i \pi/2, +i \pi/2]$  along the imaginary axis.

### Related Information

The “**casin**, **casinf**, or **casinl** Subroutine” on page 131.

---

## catan, catanf, or catanl Subroutine

### Purpose

Computes the complex arc tangent.

## Syntax

```
#include <complex.h>
```

```
double complex catan (z)
double complex z;
```

```
float complex catanf (z)
float complex z;
```

```
long double complex catanl (z)
long double complex z;
```

## Description

The **catan**, **catanf**, and **catanl** subroutines compute the complex arc tangent of  $z$ , with branch cuts outside the interval  $[-i, +i]$  along the imaginary axis.

## Parameters

$z$  Specifies the value to be computed.

## Return Values

The **catan**, **catanf**, and **catanl** subroutines return the complex arc tangent value, in the range of a strip mathematically unbounded along the imaginary axis and in the interval  $[-\pi/2, +\pi/2]$  along the real axis.

## Related Information

“catanh, catanhf, or catanhl Subroutine”

---

## catanh, catanhf, or catanhl Subroutine

### Purpose

Computes the complex arc hyperbolic tangent.

### Syntax

```
#include <complex.h>
```

```
double complex catanh (z)
double complex z;
```

```
float complex catanhf (z)
float complex z;
```

```
long double complex catanhl (z)
long double complex z;
```

### Description

The **catanh**, **catanhf**, and **catanhl** subroutines compute the complex arc hyperbolic tangent of  $z$ , with branch cuts outside the interval  $[-1, +1]$  along the real axis.

### Parameters

$z$  Specifies the value to be computed.

## Return Values

The **catanh**, **catanhf**, and **catanh1** subroutines return the complex arc hyperbolic tangent value, in the range of a strip mathematically unbounded along the real axis and in the interval  $[-i \pi/2, +i \pi/2]$  along the imaginary axis.

## Related Information

“catan, catanf, or catanl Subroutine” on page 132

---

## catclose Subroutine

### Purpose

Closes a specified message catalog.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <nl_types.h>
```

```
int catclose ( CatalogDescriptor)  
nl_catd CatalogDescriptor;
```

### Description

The **catclose** subroutine closes a specified message catalog. If your program accesses several message catalogs and you reach the maximum number of opened catalogs (specified by the **NL\_MAXOPEN** constant), you must close some catalogs before opening additional ones. If you use a file descriptor to implement the **nl\_catd** data type, the **catclose** subroutine closes that file descriptor.

The **catclose** subroutine closes a message catalog only when the number of calls it receives matches the total number of calls to the **catopen** subroutine in an application. All message buffer pointers obtained by prior calls to the **catgets** subroutine are not valid when the message catalog is closed.

### Parameters

*CatalogDescriptor*                      Points to the message catalog returned from a call to the **catopen** subroutine.

### Return Values

The **catclose** subroutine returns a value of 0 if it closes the catalog successfully, or if the number of calls it receives is fewer than the number of calls to the **catopen** subroutine.

The **catclose** subroutine returns a value of -1 if it does not succeed in closing the catalog. The **catclose** subroutine is unsuccessful if the number of calls it receives is greater than the number of calls to the **catopen** subroutine, or if the value of the *CatalogDescriptor* parameter is not valid.

## Related Information

The **catgets** (“catgets Subroutine” on page 135) subroutine, **catopen** (“catopen Subroutine” on page 136) subroutine.

For more information about the Message Facility, see Message Facility Overview for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

For more information about subroutines and libraries, see Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## catgets Subroutine

### Purpose

Retrieves a message from a catalog.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <nl_types>
```

```
char *catgets (CatalogDescriptor, SetNumber, MessageNumber, String)
nl_catd CatalogDescriptor;
int SetNumber, MessageNumber;
const char * String;
```

### Description

The **catgets** subroutine retrieves a message from a catalog after a successful call to the **catopen** subroutine. If the **catgets** subroutine finds the specified message, it loads it into an internal character string buffer, ends the message string with a null character, and returns a pointer to the buffer.

The **catgets** subroutine uses the returned pointer to reference the buffer and display the message. However, the buffer can not be referenced after the catalog is closed.

### Parameters

<i>CatalogDescriptor</i>	Specifies a catalog description that is returned by the <b>catopen</b> subroutine.
<i>SetNumber</i>	Specifies the set ID.
<i>MessageNumber</i>	Specifies the message ID. The <i>SetNumber</i> and <i>MessageNumber</i> parameters specify a particular message to retrieve in the catalog.
<i>String</i>	Specifies the default character-string buffer.

### Return Values

If the **catgets** subroutine is unsuccessful for any reason, it returns the user-supplied default message string specified by the *String* parameter.

### Related Information

The **catclose** (“catclose Subroutine” on page 134) subroutine, **catopen** (“catopen Subroutine” on page 136) subroutine.

For more information about the Message Facility, see Message Facility Overview for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

For more information about subroutines and libraries, see Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## catopen Subroutine

### Purpose

Opens a specified message catalog.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <n1_types.h>
```

```
n1_catd catopen ( CatalogName, Parameter )  
const char *CatalogName;  
int Parameter;
```

### Description

The **catopen** subroutine opens a specified message catalog and returns a catalog descriptor used to retrieve messages from the catalog. The contents of the catalog descriptor are complete when the **catgets** subroutine accesses the message catalog. The **n1\_catd** data type is used for catalog descriptors and is defined in the **n1\_types.h** file.

If the catalog file name referred to by the *CatalogName* parameter contains a leading / (slash), it is assumed to be an absolute path name. If the catalog file name is not an absolute path name, the user environment determines which directory paths to search. The **NLSPATH** environment variable defines the directory search path. When this variable is used, the **setlocale** subroutine must be called before the **catopen** subroutine.

A message catalog descriptor remains valid in a process until that process or a successful call to one of the **exec** functions closes it.

You can use two special variables, **%N** and **%L**, in the **NLSPATH** environment variable. The **%N** variable is replaced by the catalog name referred to by the call that opens the message catalog. The **%L** variable is replaced by the value of the **LC\_MESSAGES** category.

The value of the **LC\_MESSAGES** category can be set by specifying values for the **LANG**, **LC\_ALL**, or **LC\_MESSAGES** environment variable. The value of the **LC\_MESSAGES** category indicates which locale-specific directory to search for message catalogs. For example, if the **catopen** subroutine specifies a catalog with the name **mycmd**, and the environment variables are set as follows:

```
NLSPATH=../%N:../%N:/system/nls/%L/%N:/system/nls/%N LANG=fr_FR
```

then the application searches for the catalog in the following order:

```
../mycmd  
./mycmd  
/system/nls/fr_FR/mycmd  
/system/nls/mycmd
```

If you omit the **%N** variable in a directory specification within the **NLSPATH** environment variable, the application assumes that it defines a catalog name and opens it as such and will not traverse the rest of the search path.

If the **NLSPATH** environment variable is not defined, the **catopen** subroutine uses the default path. See the **/etc/environment** file for the **NLSPATH** default path. If the **LC\_MESSAGES** category is set to the



default value C, and the **LC\_\_FASTMSG** environment variable is set to true, then subsequent calls to the **catgets** subroutine generate pointers to the program-supplied default text.

The **catopen** subroutine treats the first file it finds as a message file. If you specify a non-message file in a **NLSPATH**, for example, **/usr/bin/ls**, **catopen** treats **/usr/bin/ls** as a message catalog. Thus no messages are found and default messages are returned. If you specify **/tmp** in a **NLSPATH**, **/tmp** is opened and searched for messages and default messages are displayed.

## Parameters

<i>CatalogName</i>	Specifies the catalog file to open.
<i>Parameter</i>	Determines the environment variable to use in locating the message catalog. If the value of the <i>Parameter</i> parameter is 0, use the <b>LANG</b> environment variable without regard to the <b>LC_MESSAGES</b> category to locate the catalog. If the value of the <i>Parameter</i> parameter is the <b>NL_CAT_LOCALE</b> macro, use the <b>LC_MESSAGES</b> category to locate the catalog.

## Return Values

The **catopen** subroutine returns a catalog descriptor. If the **LC\_MESSAGES** category is set to the default value C, and the **LC\_\_FASTMSG** environment variable is set to true, the **catopen** subroutine returns a value of -1.

If the **LC\_MESSAGES** category is not set to the default value C but the **catopen** subroutine returns a value of -1, an error has occurred during creation of the structure of the **nl\_catd** data type or the catalog name referred to by the *CatalogName* parameter does not exist.

## Related Information

The **catclose** (“catclose Subroutine” on page 134) subroutine, **catgets** (“catgets Subroutine” on page 135) subroutine, **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutines, **setlocale** subroutine.

The **environment** file.

For more information about the Message Facility, see the Message Facility Overview for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

For more information about subroutines and libraries, see the Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## cbrtf, cbrtl, or cbrt Subroutine

### Purpose

Computes the cube root.

### Syntax

```
#include <math.h>
```

```
float cbrtf (x)  
float x;
```

```
long double cbrtl (x)
```

```
long double x;  
  
double cbrt (x)  
double x;
```

## Description

The **cbrtf**, **cbrtl**, and **cbrt** subroutines compute the real cube root of the *x* argument.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **cbrtf**, **cbrtl**, and **cbrt** subroutines return the cube root of *x*.

If *x* is NaN, an NaN is returned.

If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* is returned.

## Related Information

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## ccos, ccosf, or ccosl Subroutine

### Purpose

Computes the complex cosine.

### Syntax

```
#include <complex.h>
```

```
double complex ccos (z)  
double complex z;
```

```
float complex ccosf (z)  
float complex z;
```

```
long double complex ccosl (z)  
long double complex z;
```

### Description

The **ccos**, **ccosf**, and **ccosl** subroutines compute the complex cosine of *z*.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **ccos**, **ccosf**, and **ccosl** subroutines return the complex cosine value.

### Related Information

“**cacos**, **cacosf**, or **cacosl** Subroutine” on page 129

---

## ccosh, ccoshf, or ccoshl Subroutine

### Purpose

Computes the complex hyperbolic cosine.

### Syntax

```
#include <complex.h>
```

```
double complex ccosh (z)
double complex z;
```

```
float complex ccoshf (z)
float complex z;
```

```
long double complex ccoshl (z)
long double complex z;
```

### Description

The **ccosh**, **ccoshf**, and **ccoshl** subroutines compute the complex hyperbolic cosine of *z*.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **ccosh**, **ccoshf**, and **ccoshl** subroutines return the complex hyperbolic cosine value.

### Related Information

“cacosh, cacoshf, or cacoshl Subroutines” on page 130

---

## ccsidtoocs or cstoccsid Subroutine

### Purpose

Provides conversion between coded character set IDs (CCSID) and code set names.

### Library

The iconv Library (**libiconv.a**)

### Syntax

```
#include <iconv.h>
```

```
CCSID cstoccsid (* Codeset)
const char *Codeset;
```

```
char *ccsidtoocs ( CCSID)
CCSID CCSID;
```

### Description

The **cstoccsid** subroutine returns the CCSID of the code set specified by the *Codeset* parameter. The **ccsidtoocs** subroutine returns the code set name of the CCSID specified by *CCSID* parameter. CCSIDs are registered IBM coded character set IDs.

## Parameters

*Codeset* Specifies the code set name to be converted to its corresponding CCSID.  
*CCSID* Specifies the CCSID to be converted to its corresponding code set name.

## Return Values

If the code set is recognized by the system, the **ctoccsid** subroutine returns the corresponding CCSID. Otherwise, null is returned.

If the CCSID is recognized by the system, the **ccsidtocs** subroutine returns the corresponding code set name. Otherwise, a null pointer is returned.

## Related Information

For more information about code set conversion, see Converters Overview for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

The National Language Support Overview for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ceil, ceilf, ceill, ceild32, ceild64, and ceild128 Subroutines

### Purpose

Compute the ceiling value.

### Syntax

```
#include <math.h>
```

```
float ceilf (x)  
float x;
```

```
long double ceill (x)  
long double x;
```

```
double ceil (x)  
double x;
```

```
_Decimal32 ceild32(x)  
_Decimal32 x;
```

```
_Decimal64 ceild64(x)  
_Decimal64 x;
```

```
_Decimal128 ceild128(x)  
_Decimal128 x;
```

### Description

The **ceilf**, **ceill**, **ceil**, **ceild32**, **ceild64**, and **ceild128** subroutines compute the smallest integral value that is not less than *x*.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the smallest integral value to be computed.

## Return Values

Upon successful completion, the **ceilf**, **ceil**, **ceil**, **ceild32**, **ceild64**, and **ceild128** subroutines return the smallest integral value that is not less than *x*, expressed as a type **float**, **long double**, **double**, **\_Decimal32**, **\_Decimal64**, or **\_Decimal128** respectively.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* is returned.

If the correct value would cause overflow, a range error occurs and the **ceilf**, **ceil**, **ceil**, **ceild32**, **ceild64**, and **ceild128** subroutines return the value of the macro **HUGE\_VALF**, **HUGE\_VALL**, **HUGE\_VAL**, **HUGE\_VAL\_D32**, **HUGE\_VAL\_D64**, and **HUGE\_VAL\_D128** respectively.

## Related Information

“**feclearexcept** Subroutine” on page 265, “**fetestexcept** Subroutine” on page 273, “**floor**, **floorf**, **floorl**, **floor32**, **floor64**, **floor128**, **nearest**, **trunc**, **itrunc**, and **utrunc** Subroutines” on page 277, and “**class**, **\_class**, **finite**, **isnan**, or **unordered** Subroutines” on page 167.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## **cexp**, **cexpf**, or **cexpl** Subroutine

### Purpose

Performs complex exponential computations.

### Syntax

```
#include <complex.h>
```

```
double complex cexp (z)
double complex z;
```

```
float complex cexpf (z)
float complex z;
```

```
long double complex cexpl (z)
long double complex z;
```

### Description

The **cexp**, **cexpf**, and **cexpl** subroutines compute the complex exponent of *z*, defined as  $e^z$ .

### Parameters

*z* Specifies the value to be computed.

## Return Values

The **cexp**, **cexpf**, and **cexpl** subroutines return the complex exponential value of *z*.

## Related Information

The “clog, clogf, or clogl Subroutine” on page 174.

---

## cfgetospeed, cfsetospeed, cfgetispeed, or cfsetispeed Subroutine

### Purpose

Gets and sets input and output baud rates.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <termios.h>
```

```
speed_t cfgetospeed ( TermiosPointer)  
const struct termios *TermiosPointer;
```

```
int cfsetospeed (TermiosPointer, Speed)  
struct termios *TermiosPointer;  
speed_t Speed;
```

```
speed_t cfgetispeed (TermiosPointer)  
const struct termios *TermiosPointer;
```

```
int cfsetispeed (TermiosPointer, Speed)  
struct termios *TermiosPointer;  
speed_t Speed;
```

### Description

The baud rate subroutines are provided for getting and setting the values of the input and output baud rates in the **termios** structure. The effects on the terminal device described below do not become effective and not all errors are detected until the **tcsetattr** function is successfully called.

The input and output baud rates are stored in the **termios** structure. The supported values for the baud rates are shown in the table that follows this discussion.

The **termios.h** file defines the type **speed\_t** as an unsigned integral type.

The **cfgetospeed** subroutine returns the output baud rate stored in the **termios** structure pointed to by the *TermiosPointer* parameter.

The **cfsetospeed** subroutine sets the output baud rate stored in the **termios** structure pointed to by the *TermiosPointer* parameter to the value specified by the *Speed* parameter.

The **cfgetispeed** subroutine returns the input baud rate stored in the **termios** structure pointed to by the *TermiosPointer* parameter.

The **cfsetispeed** subroutine sets the input baud rate stored in the **termios** structure pointed to by the *TermiosPointer* parameter to the value specified by the *Speed* parameter.

Certain values for speeds have special meanings when set in the **termios** structure and passed to the **tcsetattr** function. These values are discussed in the **tcsetattr** subroutine.

The following table lists possible baud rates:

*Baud Rate Values*

Name	Description
B0	Hang up
B5	50 baud
B75	75 baud
B110	110 baud
B134	134 baud
B150	150 baud
B200	200 baud
B300	300 baud
B600	600 baud
B1200	1200 baud
B1800	1800 baud
B2400	2400 baud
B4800	4800 baud
B9600	9600 baud
B19200	19200 baud
B38400	38400 baud

The **termios.h** file defines the name symbols of the table.

## Parameters

*TermiosPointer*                      Points to a **termios** structure.  
*Speed*                                      Specifies the baud rate.

## Return Values

The **cfgetospeed** and **cfgetispeed** subroutines return exactly the value found in the **termios** data structure, without interpretation.

Both the **cfsetospeed** and **cfsetispeed** subroutines return a value of 0 if successful and -1 if unsuccessful.

## Examples

To set the output baud rate to 0 (which forces modem control lines to stop being asserted), enter:

```
cfsetospeed (&my_termios, B0);  
tcsetattr (stdout, TCSADRAIN, &my_termios);
```

## Related Information

The **tcsetattr** subroutine.

The **termios.h** file.

## chacl or fchacl Subroutine

### Purpose

Changes the AIXC ACL type access control information of a file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/acl.h>
#include <sys/mode.h>
```

```
int chacl ( Path, ACL, ACLSize)
char *Path;
struct acl *ACL;
int ACLSize;
```

```
int fchacl ( FileDescriptor, ACL, ACLSize)
int FileDescriptor;
struct acl *ACL;
int ACLSize;
```

### Description

The **chacl** and **fchacl** subroutines set the access control attributes of a file according to the AIXC ACL Access Control List (ACL) structure pointed to by the *ACL* parameter. Note that these routines could fail if the current ACL associated with the file system object is of a different type or if the underlying physical file system does not support AIXC ACL type. It is strongly recommended that applications stop using these interfaces and instead make use of **aclx\_get/aclx\_fget** and **aclx\_put/aclx\_fput** subroutines to change the ACL.

### Parameters

*Path* Specifies the path name of the file.



ACL

Specifies the AIXC ACL to be established on the file. The format of an AIXC ACL is defined in the **sys/acl.h** file and contains the following members:

**acl\_len**

Specifies the size of the ACL (Access Control List) in bytes, including the base entries.

**Note:** The entire ACL for a file cannot exceed one memory page (4096 bytes).

**acl\_mode**

Specifies the file mode.

The following bits in the **acl\_mode** member are defined in the **sys/mode.h** file and are significant for this subroutine:

**S\_ISUID**

Enables the **setuid** attribute on an executable file.

**S\_ISGID**

Enables the **setgid** attribute on an executable file. Enables the group-inheritance attribute on a directory.

**S\_ISVTX**

Enables linking restrictions on a directory.

**S\_IXACL**

Enables extended ACL entry processing. If this attribute is not set, only the base entries (owner, group, and default) are used for access authorization checks.

Other bits in the mode, including the following, are ignored:

**u\_access**

Specifies access permissions for the file owner.

**g\_access**

Specifies access permissions for the file group.

**o\_access**

Specifies access permissions for the default class of *others*.

**acl\_ext[]**

Specifies an array of the extended entries for this access control list.

The members for the base ACL (owner, group, and others) can contain the following bits, which are defined in the **sys/access.h** file:

**R\_ACC**

Allows read permission.

**W\_ACC**

Allows write permission.

**X\_ACC** Allows execute or search permission.

*FileDescriptor*

Specifies the file descriptor of an open file.

*ACLSize*

Specifies the size of the buffer containing the ACL.

**Note:** The **chacl** subroutine requires the *Path*, *ACL*, and *ACLSize* parameters. The **fchacl** subroutine requires the *FileDescriptor*, *ACL*, and *ACLSize* parameters.

## ACL Data Structure for chacl

Each access control list structure consists of one **struct acl** structure containing one or more **struct acl\_entry** structures with one or more **struct ace\_id** structures.

If the **struct ace\_id** structure has *id\_type* set to **ACEID\_USER** or **ACEID\_GROUP**, there is only one *id\_data* element. To add multiple IDs to an ACL you must specify multiple **struct ace\_id** structures when *id\_type* is set to **ACEID\_USER** or **ACEID\_GROUP**. In this case, no error is returned for the multiple

elements, and the access checking examines only the first element. Specifically, the **errno** value **EINVAL** is not returned for *acl\_len* being incorrect in the ACL structure although more than one uid or gid is specified.

## Return Values

Upon successful completion, the **chacl** and **fchacl** subroutines return a value of 0. If the **chacl** or **fchacl** subroutine fails, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **chacl** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the disallow truncation attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>EACCESS</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.

The **chacl** or **fchacl** subroutine fails and the access control information for a file remains unchanged if one or more of the following are true:

<b>EROFS</b>	The file specified by the <i>Path</i> parameter resides on a read-only file system.
<b>EFAULT</b>	The <i>ACL</i> parameter points to a location outside of the allocated address space of the process.
<b>EINVAL</b>	The <i>ACL</i> parameter does not point to a valid ACL.
<b>EINVAL</b>	The <i>acl_len</i> member in the ACL is not valid.
<b>EIO</b>	An I/O error occurred during the operation.
<b>ENOSPC</b>	The size of the <i>ACL</i> parameter exceeds the system limit of one memory page (4KB).
<b>EPERM</b>	The effective user ID does not match the ID of the owner of the file, and the invoker does not have root user authority.

The **fchacl** subroutine fails and the file permissions remain unchanged if the following is true:

<b>EBADF</b>	The file descriptor <i>FileDescriptor</i> is not valid.
--------------	---

If Network File System (NFS) is installed on your system, the **chacl** and **fchacl** subroutines can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

## Auditing Events:

Event	Information
<b>chacl</b>	<i>Path</i>
<b>fchacl</b>	<i>FileDescriptor</i>

## Related Information

The **acl\_chg** (“acl\_chg or acl\_fchg Subroutine” on page 8) subroutine, **acl\_get** (“acl\_get or acl\_fget Subroutine” on page 10) subroutine, **acl\_put** (“acl\_put or acl\_fput Subroutine” on page 12) subroutine, **acl\_set** (“acl\_set or acl\_fset Subroutine” on page 14) subroutine, **chmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **stat** subroutine, **statacl** subroutine.

“aclx\_get or aclx\_fget Subroutine” on page 17, “aclx\_put or aclx\_fput Subroutine” on page 25.

The **aclget** command, **aclput** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## chdir Subroutine

### Purpose

Changes the current directory.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int chdir ( Path )  
const char *Path;
```

### Description

The **chdir** subroutine changes the current directory to the directory indicated by the *Path* parameter.

### Parameters

*Path*            A pointer to the path name of the directory. If the *Path* parameter refers to a symbolic link, the **chdir** subroutine sets the current directory to the directory pointed to by the symbolic link. If Network File System (NFS) is installed on the system, this path can cross into another node.

The current directory, also called the current working directory, is the starting point of searches for path names that do not begin with a / (slash). The calling process must have search access to the directory specified by the *Path* parameter.

### Return Values

Upon successful completion, the **chdir** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to identify the error.

## Error Codes

The **chdir** subroutine fails and the current directory remains unchanged if one or more of the following are true:

<b>EACCES</b>	Search access is denied for the named directory.
<b>ENOENT</b>	The named directory does not exist.
<b>ENOTDIR</b>	The path name is not a directory.

The **chdir** subroutine can also be unsuccessful for other reasons. See Appendix A, Base Operating System Error Codes for Services That Require Path-Name Resolution (Appendix A, “Base Operating System Error Codes for Services That Require Path-Name Resolution,” on page 1345) for a list of additional error codes.

If NFS is installed on the system, the **chdir** subroutine can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **chroot** (“chroot Subroutine” on page 160) subroutine.

The **cd** command.

Appendix A, “Base Operating System Error Codes for Services That Require Path-Name Resolution,” on page 1345.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## chmod or fchmod Subroutine

### Purpose

Changes file system object’s base file mode bits.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/stat.h>
```

```
int chmod ( Path, Mode)
const char *Path;
mode_t Mode;
```

```
int fchmod ( FileDescriptor, Mode)
int FileDescriptor;
mode_t Mode;
```

### Description

The **chmod** subroutine sets the access permissions of the file specified by the *Path* parameter. If Network File System (NFS) is installed on your system, this path can cross into another node.

Use the **fchmod** subroutine to set the access permissions of an open file pointed to by the *FileDescriptor* parameter.

If *FileDescriptor* references a shared memory object, the **fchmod** subroutine affects the **S\_IRUSR**, **S\_IWUSR**, **S\_IRGRP**, **S\_IWGRP**, **S\_IROTH**, and **S\_IWOTH** file permission bits.

The access control information is set according to the *Mode* parameter. Note that these routines will replace any existing ACL associated with the file system object.

## Parameters

*FileDescriptor*                      Specifies the file descriptor of an open file or shared memory object.

*Mode*

Specifies the bit pattern that determines the access permissions. The *Mode* parameter is constructed by logically ORing one or more of the following values, which are defined in the **sys/mode.h** file:

**S\_ISUID**

Enables the **setuid** attribute for an executable file. A process executing this program acquires the access rights of the owner of the file.

**S\_ISGID**

Enables the **setgid** attribute for an executable file. A process executing this program acquires the access rights of the group of the file. Also, enables the group-inheritance attribute for a directory. Files created in this directory have a group equal to the group of the directory.

The following attributes apply only to files that are directly executable. They have no meaning when applied to executable text files such as shell scripts and **awk** scripts.

**S\_ISVTX**

Enables the **link/unlink** attribute for a directory. Files cannot be linked to in this directory. Files can only be unlinked if the requesting process has write permission for the directory and is either the owner of the file or the directory.

**S\_ISVTX**

Enables the **save text** attribute for an executable file. The program is not unmapped after usage.

**S\_ENFMT**

Enables enforcement-mode record locking for a regular file. File locks requested with the **lockf** subroutine are enforced.

**S\_IRUSR**

Permits the file's owner to read it.

**S\_IWUSR**

Permits the file's owner to write to it.

**S\_IXUSR**

Permits the file's owner to execute it (or to search the directory).

**S\_IRGRP**

Permits the file's group to read it.

**S\_IWGRP**

Permits the file's group to write to it.

**S\_IXGRP**

Permits the file's group to execute it (or to search the directory).

**S\_IROTH**

Permits others to read the file.

**S\_IWOTH**

Permits others to write to the file.

**S\_IXOTH**

Permits others to execute the file (or to search the directory).

Other mode values exist that can be set with the **mknod** subroutine but not with the **chmod** subroutine.

*Path*

Specifies the full path name of the file.

## Return Values

Upon successful completion, the **chmod** subroutine and **fchmod** subroutines return a value of 0. If the **chmod** subroutine or **fchmod** subroutine is unsuccessful, a value of -1 is returned, and the **errno** global variable is set to identify the error.

## Error Codes

The **chmod** subroutine is unsuccessful and the file permissions remain unchanged if one of the following is true:

<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENOENT</b>	The named file does not exist.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.

The **fchmod** subroutine is unsuccessful and the file permissions remain unchanged if the following is true:

<b>EBADF</b>	The value of the <i>FileDescriptor</i> parameter is not valid.
--------------	--

The **chmod** or **fchmod** subroutine is unsuccessful and the access control information for a file remains unchanged if one of the following is true:

<b>EPERM</b>	The effective user ID does not match the owner of the file, and the process does not have appropriate privileges.
<b>EROFS</b>	The named file resides on a read-only file system.
<b>EIO</b>	An I/O error occurred during the operation.

If NFS is installed on your system, the **chmod** and **fchmod** subroutines can also be unsuccessful if the following is true:

<b>ESTALE</b>	The root or current directory of the process is located in a virtual file system that has been unmounted.
<b>ETIMEDOUT</b>	The connection timed out.

## Security

Access Control: The invoker must have search permission for all components of the *Path* prefix.

If you receive the **EBUSY** error, toggle the **enforced locking** attribute in the *Mode* parameter and retry your operation. The **enforced locking** attribute should never be used on a file that is part of the Trusted Computing Base.

## Related Information

The **acl\_chg** (“acl\_chg or acl\_fchg Subroutine” on page 8) subroutine, **acl\_get** (“acl\_get or acl\_fget Subroutine” on page 10) subroutine, **acl\_put** (“acl\_put or acl\_fput Subroutine” on page 12) subroutine, **acl\_set** (“acl\_set or acl\_fset Subroutine” on page 14) subroutine, **chacl** (“chacl or fchacl Subroutine” on page 144) subroutine, **statacl** subroutine, **stat** subroutine.

“aclx\_get or aclx\_fget Subroutine” on page 17, “aclx\_put or aclx\_fput Subroutine” on page 25.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## chown, fchown, lchown, chownx, or fchownx Subroutine

### Purpose

Changes file ownership.

### Library

Standard C Library (**libc.a**)

## Syntax

Syntax for the **chown**, **fchown**, and **lchown** Subroutines:

```
#include <sys/types.h>
```

```
#include <unistd.h>
```

```
int chown ( Path, Owner, Group)
```

```
const char *Path;
```

```
uid_t Owner;
```

```
gid_t Group;
```

```
int fchown ( FileDescriptor, Owner, Group)
```

```
int FileDescriptor;
```

```
uid_t Owner;
```

```
gid_t Group;
```

```
int lchown ( Path, Owner, Group)
```

```
const char *fname
```

```
uid_t uid
```

```
gid_t gid
```

Syntax for the **chownx** and **fchownx** Subroutines:

```
#include <sys/types.h>
```

```
#include <sys/chownx.h>
```

```
int chownx (Path, Owner, Group, Flags)
```

```
char *Path;
```

```
uid_t Owner;
```

```
gid_t Group;
```

```
int Flags;
```

```
int fchownx (FileDescriptor, Owner, Group, Flags)
```

```
int FileDescriptor;
```

```
uid_t Owner;
```

```
gid_t Group;
```

```
int Flags;
```

## Description

The **chown**, **chownx**, **fchown**, **fchownx**, and **lchown** subroutines set the file owner and group IDs of the specified file system object. Root user authority is required to change the owner of a file.

A function **lchown** function sets the owner ID and group ID of the named file similarly to **chown** function except in the case where the named file is a symbolic link. In this case **lchown** function changes the ownership of the symbolic link file itself, while **chown** function changes the ownership of the file or directory to which the symbolic link refers.

## Parameters

*FileDescriptor* Specifies the file descriptor of an open file.



<i>Flags</i>	Specifies whether the file owner ID or group ID should be changed. This parameter is constructed by logically ORing the following values: <b>T_OWNER_AS_IS</b> Ignores the value specified by the <i>Owner</i> parameter and leaves the owner ID of the file unaltered. <b>T_GROUP_AS_IS</b> Ignores the value specified by the <i>Group</i> parameter and leaves the group ID of the file unaltered.
<i>Group</i>	Specifies the new group of the file. For the <b>chown</b> , <b>fchown</b> , and <b>lchown</b> commands, if this value is -1, the group is not changed. (A value of -1 indicates only that the group is not changed; it does not indicate a group that is not valid. An owner or group ID cannot be invalid.) For the <b>chownx</b> and <b>fchownx</b> commands, the subroutines change the Group to -1 if -1 is supplied for Group and <b>T_GROUP_AS_IS</b> is not set.
<i>Owner</i>	Specifies the new owner of the file. For the <b>chown</b> , <b>fchown</b> , and <b>lchown</b> commands, if this value is -1, the group is not changed. (A value of -1 indicates only that the group is not changed; it does not indicate a group that is not valid. An owner or group ID cannot be invalid.) For the <b>chownx</b> and <b>fchownx</b> commands, the subroutines change the Owner to -1 if -1 is supplied for Owner and <b>T_OWNER_AS_IS</b> is not set.
<i>Path</i>	Specifies the full path name of the file. If <i>Path</i> resolves to a symbolic link, the ownership of the file or directory pointed to by the symbolic link is changed.

## Return Values

Upon successful completion, the **chown**, **chownx**, **fchown**, **fchownx**, and **lchown** subroutines return a value of 0. If the **chown**, **chownx**, **fchown**, **fchownx**, or **lchown** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **chown**, **chownx**, or **lchown** subroutine is unsuccessful and the owner and group of a file remain unchanged if one of the following is true:

<b>EACCES</b>	Search permission is denied on a component of the <i>Path</i> parameter.
<b>EDQUOT</b>	The new group for the file system object cannot be set because the group's quota of disk blocks or i-nodes has been exhausted on the file system.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.
<b>EINVAL</b>	The owner or group ID supplied is not valid.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist; or a component of the <i>Path</i> parameter does not exist; or the process has the <b>disallow truncation</b> attribute set; or the <i>Path</i> parameter is null.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>EPERM</b>	The effective user ID does not match the owner of the file, and the calling process does not have the appropriate privileges.
<b>EROFS</b>	The named file resides on a read-only file system.
<b>ESTALE</b>	The root or current directory of the process is located in a virtual file system that has been unmounted.

The **fchown** or **fchownx** subroutine is unsuccessful and the file owner and group remain unchanged if one of the following is true:

<b>EBADF</b>	The named file resides on a read-only file system.
--------------	--

<b>EDQUOT</b>	The new group for the file system object cannot be set because the group's quota of disk blocks or i-nodes has been exhausted on the file system.
<b>EIO</b>	An I/O error occurred during the operation.

## Security

Access Control: The invoker must have search permission for all components of the *Path* parameter.

---

## chpass Subroutine

### Purpose

Changes user passwords.

### Library

Standard C Library (**libc.a**)

Thread Safe Security Library (**libs\_r.a**)

### Syntax

```
int chpass (UserName, Response, Reenter, Message)
char *UserName;
char *Response;
int *Reenter;
char **Message;
```

### Description

The **chpass** subroutine maintains the requirements that the user must meet to change a password. This subroutine is the basic building block for changing passwords and handles password changes for local, NIS, and DCE user passwords.

The *Message* parameter provides a series of messages asking for old and new passwords, or providing informational messages, such as the reason for a password change failing. The first *Message* prompt is a prompt for the old password. This parameter does not prompt for the old password if the user has a real user ID of 0 (zero) and is changing a local user, or if the user has no current password. The **chpass** subroutine does not prompt a user with root authority for an old password. It informs the program that no message was sent and that it should invoke **chpass** again. If the user satisfies the first *Message* parameter's prompt, the system prompts the user to enter the new password. Each message is contained in the *Message* parameter and is displayed to the user. The *Response* parameter returns the user's response to the **chpass** subroutine.

The *Reenter* parameter indicates when a user has satisfied all prompt messages. The parameter remains nonzero until a user has passed all prompts. After the returned value of *Reenter* is 0, the return code signals whether the password change has succeeded or failed. When progressing through prompts for a user, the value of *Reenter* must be maintained by the caller between invocations of **chpass**.

The **chpass** subroutine maintains internal state information concerning the next prompt message to present to the user. If the calling program supplies a different user name before all prompt messages are complete for the user, the internal state information is reset and prompt messages begin again. State information is also kept in the *Reenter* variable. The calling program must maintain the value of *Reenter* between calls to **chpass**.

The **chpass** subroutine determines the administration domain to use during password changes. It determines if the user is defined locally, defined in Network Information Service (NIS), or defined in

Distributed Computing Environment (DCE). Password changes occur only in these domains. System administrators may override this convention with the registry value in the `/etc/security/user` file. If the registry value is defined, the password change can only occur in the specified domain. System administrators can use this registry value if the user is administered on a remote machine that periodically goes down. If the user is allowed to log in through some other authentication method while the server is down, password changes remain to follow only the primary server.

The **chpass** subroutine allows the user to change passwords in two ways. For normal (non-administrative) password changes, the user must supply the old password, either on the first call to the **chpass** subroutine or in response to the first message from **chpass**. If the user is root, real user ID of 0, local administrative password changes are handled by supplying a null pointer for the *Response* parameter during the initial call

Users that are not administered locally are always queried for their old password.

The **chpass** subroutine is always in one of the following states:

1. Initial state: The caller invokes the **chpass** subroutine with NULL *response* parameter and receives the initial password prompt in the *message* parameter.
2. Verify initial password: The caller invokes the **chpass** subroutine with the results of prompting the user with earlier *message* parameter as the *response* parameter. The caller is given a prompt to enter the new password in the *message* parameter.
3. Enter new password: The caller invokes the **chpass** subroutine with the results of prompting user with the new password prompt in the *response* parameter. The caller will be given a prompt to repeat the new password in the *message* parameter.
4. Verify new password: The caller invokes the **chpass** subroutine with the results of prompting the user to repeat the new password in the *response* parameter. The **chpass** subroutine then performs the actual password change.

Any step in the above process can result in the **chpass** subroutine terminating the dialog. This is signalled when the *reenter* variable is set to 0. The return code indicates the nature of the failure.

**Note:** Set the *setuid* and owner to root for your own programs that use the **chpass** subroutine.

## Parameters

<i>UserName</i>	Specifies the user's name whose password is to be changed.
<i>Response</i>	Specifies a character string containing the user's response to the last prompt.
<i>Reenter</i>	Points to a Boolean value used to signal whether the <b>chpass</b> subroutine has completed processing. If the <i>Reenter</i> parameter is a nonzero value, the <b>chpass</b> subroutine expects the user to satisfy the prompt message provided by the <i>Message</i> parameter. If the <i>Reenter</i> parameter is 0, the <b>chpass</b> subroutine has completed processing.
<i>Message</i>	Points to a pointer that the <b>chpass</b> subroutine allocates memory for and fills in. This replacement string is then suitable for printing and issues challenge messages (if the <i>Reenter</i> parameter is a nonzero value). The string can also issue informational messages such as why the user failed to change the password (if the <i>Reenter</i> parameter is 0). The calling application is responsible for freeing this memory.

## Return Values

Upon successful completion, the **chpass** subroutine returns a value of 0. If the **chpass** subroutine is unsuccessful, it returns the following values:

- 1 Indicates the call failed in the thread safe library **libs\_r.a**. **ERRNO** will indicate the failure code.
- 1 Indicates that the password change was unsuccessful and the user should attempt again. This return value occurs if a password restriction is not met, such as if the password is not long enough.

- 2 Indicates that the password change was unsuccessful and the user should not attempt again. This return value occurs if the user enters an incorrect old password or if the network is down (the password change cannot occur).

## Error Codes

The **chpass** subroutine is unsuccessful if one of the following values is true:

<b>ENOENT</b>	Indicates that the user cannot be found.
<b>ESAD</b>	Indicates that the user did not meet the criteria to change the password.
<b>EPERM</b>	Indicates that the user did not have permission to change the password.
<b>EINVAL</b>	Indicates that the parameters are not valid.
<b>ENOMEM</b>	Indicates that memory allocation (malloc) failed.

## Related Information

The **authenticate** (“authenticate Subroutine” on page 113) subroutine.

---

## chpassx Subroutine

### Purpose

Changes multiple method passwords.

### Library

Standard C Library (**libc.a**)

Thread Safe Security Library (**libs\_r.a**)

### Syntax

```
int chpassx (UserName, Response, Reenter, Message, State)
char *UserName;
char *Response;
int *Reenter;
char **Message;
void **State;
```

### Description

The **chpassx** subroutine maintains the requirements that the user must meet to change a password. This subroutine is the basic building block for changing passwords, and it handles password changes for local, NIS, and loadable authentication module user passwords. It uses information provided by the **authenticatex** and **passwdexpiredx** subroutines to indicate which passwords were used when a user authenticated and whether or not those passwords are expired.

The *Message* parameter provides a series of messages asking for old and new passwords, or providing informational messages, such as the reason for a password change failing. The first *Message* prompt is a prompt for the old password. This parameter does not prompt for the old password if the user has a real user ID of 0 and is changing a local user, or if the user has no current password. The **chpassx** subroutine does not prompt a user with root authority for an old password when only a local password is being changed. It informs the program that no message was sent and that it should invoke **chpass** again. If the user satisfies the first *Message* parameter’s prompt, the system prompts the user to enter the new password. Each message is contained in the *Message* parameter and is displayed to the user. The *Response* parameter returns the user’s response to the **chpassx** subroutine.

The *Reenter* parameter remains a nonzero value until the user satisfies all of the prompt messages or until the user incorrectly responds to a prompt message. When the *Reenter* parameter is 0, the return code signals whether the password change completed or failed. The calling application must initialize the *Reenter* parameter to 0 before the first call to the **chpassx** subroutine and the application cannot modify the *Reenter* parameter until the sequence of **chpassx** subroutine calls has completed.

The **authenticatex** subroutine ascertains the authentication domains the user can attempt. The subroutine uses the **SYSTEM** attribute for the user. Each token that is displayed in the **SYSTEM** line corresponds to a method that can be dynamically loaded and processed. Likewise, the system can provide multiple or alternate authentication paths.

The *State* parameter contains information from an earlier call to the **authenticatex** or **passwdexpirex** subroutines. That information indicates which administration domains were used when the user was authenticated and which passwords have expired and can be changed by the user. The *State* parameter must be initialized to null when the **chpassx** subroutine is not being called after an earlier call to the **authenticatex** or **passwdexpiredx** subroutines, or if the calling program does not wish to use the information from an earlier call.

The **chpassx** subroutine maintains internal state information concerning the next prompt message to present to the user. If the calling program supplies a different user name before all prompt messages are complete for the user, the internal state information is reset and prompt messages begin again.

The **chpassx** subroutine determines the administration domain to use during password changes. It determines if the user is defined locally, defined in Network Information Service (NIS), defined in Distributed Computing Environment (DCE), or defined in another administrative domain supported by a loadable authentication module. Password changes use the user's **SYSTEM** attribute and information in the *State* parameter. When the *State* parameter includes information from an earlier call to the **authenticatex** subroutine, only the administrative domains that were used for authentication are changed. When the *State* parameter includes information from an earlier call to the **passwdexpiredx** subroutine, only the administrative domains that have expired passwords are changed. The *State* parameter can contain information from calls to both **authenticatex** and **passwdexpiredx**, in which case passwords that were used for authentication are changed, even if they are not expired, so that passwords remain synchronized between administrative domains.

The **chpassx** subroutine allows the user to change passwords in two ways. For normal (nonadministrative) password changes, the user must supply the old password, either on the first call to the **chpassx** subroutine or in response to the first message from **chpassx**. If the user is root (with a real user ID of 0), local administrative password changes are handled by supplying a null pointer for the *Response* parameter during the initial call.

Users that are not administered locally are always queried for their old password.

The **chpassx** subroutine is always in one of three states: entering the old password, entering the new password, or entering the new password again. If any of these states do not need to be complied with, the **chpassx** subroutine returns a null challenge.

## Parameters

*Message* Points to a pointer that the **chpassx** subroutine allocates memory for and fills in. This replacement string is then suitable for printing and issues challenge messages (if the *Reenter* parameter is a nonzero value). The string can also issue informational messages, such as why the user failed to change the password (if the *Reenter* parameter is 0). The calling application is responsible for freeing this memory.

<i>Reenter</i>	Points to an integer value used to signal whether the <b>chpassx</b> subroutine has completed processing. If the <i>Reenter</i> parameter is a nonzero value, the <b>chpassx</b> subroutine expects the user to satisfy the prompt message provided by the <i>Message</i> parameter. If the <i>Reenter</i> parameter is 0, the <b>chpassx</b> subroutine has completed processing.
<i>Response</i>	Specifies a character string containing the user's response to the last prompt.
<i>State</i>	Points to a pointer that the <b>chpassx</b> subroutine allocates memory for and fills in. The <i>State</i> parameter can also be the result of an earlier call to the <b>authenticatex</b> or <b>passwdexpiredx</b> subroutines. This parameter contains information about each password that has been changed for the user. The calling application is responsible for freeing this memory after the <b>chpassx</b> subroutine has completed.
<i>UserName</i>	Specifies the user's name whose password is to be changed.

## Return Values

Upon successful completion, the **chpassx** subroutine returns a value of 0. If this subroutine fails, it returns the following values:

- 1 The call failed in the **libs\_r.a** thread safe library. **errno** indicates the failure code.
- 1 The password change was unsuccessful and the user should try again. This return value occurs if a password restriction is not met (for example, the password is not long enough).
- 2 The password change was unsuccessful and the user should not try again. This return value occurs if the user enters an incorrect old password or if the network is down (the password change cannot occur).

## Error Codes

The **chpassx** subroutine is unsuccessful if one of the following values is true:

<b>EINVAL</b>	The parameters are not valid.
<b>ENOENT</b>	The user cannot be found.
<b>ENOMEM</b>	Memory allocation (malloc) failed.
<b>EPERM</b>	The user did not have permission to change the password.
<b>ESAD</b>	The user did not meet the criteria to change the password.

## Related Information

The “authenticatex Subroutine” on page 115, “passwdexpiredx Subroutine” on page 974.

---

## chprojattr Subroutine

### Purpose

Updates and modifies the project attributes in kernel project registry for the given project.

### Library

The **libaacct.a** library.

### Syntax

```
<sys/aacct.h>
```

```
chprojattr(struct project *, int cmd)
```

### Description

The **chprojattr** subroutine alters the attributes of a project defined in the kernel project registry. A pointer to struct project containing the project definition and the operation command is sent as input arguments. The following operations are permitted:

- PROJ\_ENABLE\_AGGR - Enables aggregation for the specified project
- PROJ\_DISABLE\_AGGR - Disables aggregation for the specified project

If PROJ\_ENABLE\_AGGR is passed, then the aggregation status bit is set to 1. If PROJ\_DISABLE\_AGGR is passed, then the aggregation status bit set to 0.

**Note:** To initialize the project structure, the user must call the **getprojdef** subroutine before calling the **chprojattr** subroutine.

## Parameters

<i>project</i>	Pointer containing the project definition.
<i>cmd</i>	An integer command indicating whether to perform a set or clear operation.

## Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid arguments passed. The passed command flag is invalid or the passed pointer is NULL.
<b>ENONENT</b>	Project not found.

## Related Information

The “addproj Subroutine” on page 31, “chprojattrdb Subroutine,” “getproj Subroutine” on page 418, “getprojs Subroutine” on page 420, rmproj Subroutine.

---

## chprojattrdb Subroutine

### Purpose

Updates the project attributes in the project database.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
chprojattrdb(void *handle, struct project *project, int cmd)
```

### Description

The **chprojattrdb** subroutine alters the attributes of the named project in the specified project database, which is controlled through the *handle* parameter. The following commands are permitted:

- **PROJ\_ENABLE\_AGGR** — Enables aggregation for the specified project

- **PROJ\_DISABLE\_AGGR** — Disables aggregation for the specified project

The project database must be initialized before calling this subroutine. The **projdballoc** subroutine is provided for this purpose. The **chprojattrdb** subroutine must be called after the **getprojdb** subroutine, which sets the record pointer to point to the project that needs to be modified.

**Note:** The **chprojattrdb** subroutine must be called after the **getprojdb** subroutine, which makes the named project the current project.

## Parameters

<i>handle</i>	Pointer to the handle allocated for the project database.
<i>project</i>	Pointer containing the project definition.
<i>cmd</i>	An integer command indicating whether to perform a set or clear operation.

## Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid arguments passed. The passed command flag is invalid or the passed pointer is NULL.
<b>ENONENT</b>	Project not found.

## Related Information

The “addprojdb Subroutine” on page 32, “chprojattr Subroutine” on page 158, “getfirstprojdb Subroutine” on page 368, “getnextprojdb Subroutine” on page 396, “getprojdb Subroutine” on page 419, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, “projdbfree Subroutine” on page 1181, rmpojdb Subroutine.

---

## chroot Subroutine

### Purpose

Changes the effective root directory.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int chroot (const char * Path)
char *Path;
```



## Description

The **chroot** subroutine causes the directory named by the *Path* parameter to become the effective root directory. If the *Path* parameter refers to a symbolic link, the **chroot** subroutine sets the effective root directory to the directory pointed to by the symbolic link. If Network File System (NFS) is installed on your system, this path can cross into another node.

The effective root directory is the starting point when searching for a file's path name that begins with / (slash). The current directory is not affected by the **chroot** subroutine.

The calling process must have root user authority in order to change the effective root directory. The calling process must also have search access to the new effective root directory.

The .. (double period) entry in the effective root directory is interpreted to mean the effective root directory itself. Thus, this directory cannot be used to access files outside the subtree rooted at the effective root directory.

## Parameters

*Path*            Pointer to the new effective root directory.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **chroot** subroutine fails and the effective root directory remains unchanged if one or more of the following are true:

<b>ENOENT</b>	The named directory does not exist.
<b>EACCES</b>	The named directory denies search access.
<b>EPERM</b>	The process does not have root user authority.

The **chroot** subroutine can be unsuccessful for other reasons. See Appendix A. Base Operating System Error Codes for Services that Require Path-Name Resolution (Appendix A, "Base Operating System Error Codes for Services That Require Path-Name Resolution," on page 1345) for a list of additional errors.

If NFS is installed on the system, the **chroot** subroutine can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **chdir** ("chdir Subroutine" on page 147) subroutine.

The **chroot** command.

Appendix A, "Base Operating System Error Codes for Services That Require Path-Name Resolution," on page 1345.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## chssys Subroutine

### Purpose

Modifies the subsystem objects associated with the *SubsystemName* parameter.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <sys/srcobj.h>
#include <spc.h>
```

```
int chssys( SubsystemName, SRCSubsystem)
char *SubsystemName;
struct SRCsubsys *SRCSubsystem;
```

### Description

The **chssys** subroutine modifies the subsystem objects associated with the specified subsystem with the values in the **SRCsubsys** structure. This action modifies the objects associated with subsystem in the following object classes:

- Subsystem Environment
- Subserver Type
- Notify

The Subserver Type and Notify object classes are updated only if the subsystem name has been changed.

The **SRCsubsys** structure is defined in the */usr/include/sys/srcobj.h* file.

The program running with this subroutine must be running with the group system.

### Parameters

<i>SRCSubsystem</i>	Points to the <b>SRCsubsys</b> structure.
<i>SubsystemName</i>	Specifies the name of the subsystem.

### Return Values

Upon successful completion, the **chssys** subroutine returns a value of 0. Otherwise, it returns a value of -1 and the **odmerrno** variable is set to indicate the error, or a System Resource Controller (SRC) error code is returned.

### Error Codes

The **chssys** subroutine is unsuccessful if one or more of the following are true:

<b>SRC_NONAME</b>	No subsystem name is specified.
<b>SRC_NOPATH</b>	No subsystem path is specified.
<b>SRC_BADNSIG</b>	Invalid stop normal signal.
<b>SRC_BADFSIG</b>	Invalid stop force signal.
<b>SRC_NOCONTACT</b>	Contact not signal, sockets, or message queues.
<b>SRC_SSME</b>	Subsystem name does not exist.
<b>SRC_SUBEXIST</b>	New subsystem name is already on file.
<b>SRC_SYNEXIST</b>	New subsystem synonym name is already on file.

<b>SRC_NOREC</b>	The specified <b>SRCsubsys</b> record does not exist.
<b>SRC_SUBSYS2BIG</b>	Subsystem name is too long.
<b>SRC_SYN2BIG</b>	Synonym name is too long.
<b>SRC_CMDARG2BIG</b>	Command arguments are too long.
<b>SRC_PATH2BIG</b>	Subsystem path is too long.
<b>SRC_STDIN2BIG</b>	stdin path is too long.
<b>SRC_STDOUT2BIG</b>	stdout path is too long.
<b>SRC_STDERR2BIG</b>	stderr path is too long.
<b>SRC_GRPNAM2BIG</b>	Group name is too long.

## Security

Privilege Control: This command has the Trusted Path attribute. It has the following kernel privilege:

**SET\_PROC\_AUDIT** kernel privilege

Files Accessed:

Mode	File
644	<b>/etc/objrepos/SRCsubsys</b>
644	<b>/etc/objrepos/SRCsubsvr</b>
644	<b>/etc/objrepos/SRCnotify</b>

Auditing Events:

Event	Information
<b>SRC_Chssys</b>	

## Files

<b>/etc/objrepos/SRCsubsys</b>	SRC Subsystem Configuration object class.
<b>/etc/objrepos/SRCsubsvr</b>	SRC Subserver Configuration object class.
<b>/etc/objrepos/SRCnotify</b>	SRC Notify Method object class.
<b>/dev/SRC</b>	Specifies the <b>AF_UNIX</b> socket file.
<b>/dev/.SRC-unix</b>	Specifies the location for temporary socket files.

## Related Information

The **addssys** (“addssys Subroutine” on page 33) subroutine, **delssys** (“delssys Subroutine” on page 212) subroutine.

The **chssys** command, **mkssys** command, **rmssys** command.

System Resource Controller in *Operating system and device management*.

Defining Your Subsystem to the SRC, List of SRC Subroutines, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## cimag, cimagn, or cimagn Subroutine

### Purpose

Performs complex imaginary computations.

## Syntax

```
#include <complex.h>
```

```
double cimag (z)  
double complex z;
```

```
float cimagf (z)  
float complex z;
```

```
long double cimagl (z)  
long double complex z;
```

## Description

The **cimag**, **cimagf**, and **cimagl** subroutines compute the imaginary part of *z*.

## Parameters

*z* Specifies the value to be computed.

## Return Values

The **cimag**, **cimagf**, and **cimagl** subroutines return the imaginary part value (as a real).

## Related Information

“**carg**, **cargf**, or **cargl** Subroutine” on page 131, “**conj**, **conjf**, or **conjl** Subroutine” on page 182, “**cproj**, **cprojf**, or **cprojl** Subroutine” on page 190, and “**creal**, **crealf**, or **creall** Subroutine” on page 190.

---

## ckuseracct Subroutine

### Purpose

Checks the validity of a user account.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <login.h>
```

```
int ckuseracct ( Name, Mode, TTY)  
char *Name;  
int Mode;  
char *TTY;
```

### Description

**Note:** This subroutine is obsolete and is provided only for backwards compatibility. Use the **loginrestrictions** subroutine, which performs a superset of the functions of the **ckuseracct** subroutine, instead.

The **ckuseracct** subroutine checks the validity of the user account specified by the *Name* parameter. The *Mode* parameter gives the mode of the account usage, and the *TTY* parameter defines the terminal being used for the access. The **ckuseracct** subroutine checks for the following conditions:

- Account existence

- Account expiration

The *Mode* parameter specifies other mode-specific checks.

## Parameters

<i>Name</i>	Specifies the login name of the user whose account is to be validated.
<i>Mode</i>	Specifies the manner of usage. Valid values as defined in the <b>login.h</b> file are listed below. The <i>Mode</i> parameter must be one of these or 0:
<b>S_LOGIN</b>	Verifies that local logins are permitted for this account.
<b>S_SU</b>	Verifies that the <b>su</b> command is permitted and that the current process has a group ID that can invoke the <b>su</b> command to switch to the account.
<b>S_DAEMON</b>	Verifies the account can be used to invoke daemon or batch programs using the <b>src</b> or <b>cron</b> subsystems.
<b>S_RLOGIN</b>	Verifies the account can be used for remote logins using the <b>rlogind</b> or <b>telnetd</b> programs.
<i>TTY</i>	Specifies the terminal of the originating activity. If this parameter is a null pointer or a null string, no TTY origin checking is done.

## Security

Files Accessed:

Mode	File
r	/etc/passwd
r	/etc/security/user

## Return Values

If the account is valid for the specified usage, the **ckuseracct** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to the appropriate error code.

## Error Codes

The **ckuseracct** subroutine fails if one or more of the following are true:

<b>ENOENT</b>	The user specified in the <i>Name</i> parameter does not have an account.
<b>ESTALE</b>	The user's account is expired.
<b>EACCES</b>	The specified terminal does not have access to the specified account.
<b>EACCES</b>	The <i>Mode</i> parameter is <b>S_SU</b> , and the current process is not permitted to use the <b>su</b> command to access the specified user.
<b>EACCES</b>	Access to the account is not permitted in the specified <i>Mode</i> .
<b>EINVAL</b>	The <i>Mode</i> parameter is not one of <b>S_LOGIN</b> , <b>S_SU</b> , <b>S_DAEMON</b> , <b>S_RLOGIN</b> .

## Related Information

The **ckuserID** (“ckuserID Subroutine” on page 166) subroutine, **getpcred** (“getpcred Subroutine” on page 403) subroutine, **getpenv** (“getpenv Subroutine” on page 405) subroutine, **setpcred** subroutine, **setpenv** subroutine.

The **login** command, **rlogin** command, **su** command, **telnet** command.

The **cron** daemon.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ckuserID Subroutine

### Purpose

Authenticates the user.

**Note:** This subroutine is obsolete and is provided for backwards compatibility. Use the **authenticate** (“authenticate Subroutine” on page 113) subroutine, instead.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <login.h>
int ckuserID ( User, Mode)
int Mode;
char *User;
```

### Description

The **ckuserID** subroutine authenticates the account specified by the *User* parameter. The mode of the authentication is given by the *Mode* parameter. The **login** and **su** commands continue to use the **ckuserID** subroutine to process the **/etc/security/user auth1** and **auth2** authentication methods.

The **ckuserID** subroutine depends on the **authenticate** (“authenticate Subroutine” on page 113) subroutine to process the **SYSTEM** attribute in the **/etc/security/user** file. If authentication is successful, the **passwdexpired** (“passwdexpired Subroutine” on page 973) subroutine is called.

Errors caused by grammar or load modules during a call to the **authenticate** subroutine are displayed to the user if the user was authenticated. These errors are audited with the **USER\_Login** audit event if the user failed authentication.

### Parameters

*User* Specifies the name of the user to be authenticated.  
*Mode* Specifies the mode of authentication. This parameter is a bit mask and may contain one or more of the following values, which are defined in the **login.h** file:

#### **S\_PRIMARY**

The primary authentication methods defined for the *User* parameter are checked. All primary authentication checks must be passed.

#### **S\_SECONDARY**

The secondary authentication methods defined for the *User* parameter are checked. Secondary authentication checks are not required to be successful.

Primary and secondary authentication methods for each user are set in the **/etc/security/user** file by defining the **auth1** and **auth2** attributes. If no primary methods are defined for a user, the **SYSTEM** attribute is assumed. If no secondary methods are defined, there is no default.

## Security

Files Accessed:

Mode	File
r	/etc/passwd
r	/etc/security/passwd
r	/etc/security/user
r	/etc/security/login.cfg

## Return Values

If the account is valid for the specified usage, the **ckuserID** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **ckuserID** subroutine fails if one or more of the following are true:

<b>ESAD</b>	Security authentication failed for the user.
<b>EINVAL</b>	The <i>Mode</i> parameter is neither <b>S_PRIMARY</b> nor <b>S_SECONDARY</b> or the <i>Mode</i> parameter is both <b>S_PRIMARY</b> and <b>S_SECONDARY</b> .

## Related Information

The **authenticate** (“authenticate Subroutine” on page 113) subroutine, **ckuseracct** (“ckuseracct Subroutine” on page 164) subroutine, **getpcred** (“getpcred Subroutine” on page 403) subroutine, **getpenv** (“getpenv Subroutine” on page 405) subroutine, **passwdexpired** (“passwdexpired Subroutine” on page 973) subroutine, **setpcred** subroutine, **setpenv** subroutine.

The **login** command, **su** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## class, \_class, finite, isnan, or unordered Subroutines

### Purpose

Determines classifications of floating-point numbers.

### Libraries

IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>  
#include <float.h>
```

```
int  
class( x)  
double x;
```

```
#include <math.h>  
#include <float.h>
```

```

int
_class( x)
double x;
#include <math.h>
int finite(x)
double x;
#include <math.h>
int isnan(x)
double x;
#include <math.h>

int unordered(x, y)
double x, y;

```

## Description

The **class** subroutine, **\_class** subroutine, **finite** subroutine, **isnan** subroutine, and **unordered** subroutine determine the classification of their floating-point value. The **unordered** subroutine determines if a floating-point comparison involving *x* and *y* would generate the IEEE floating-point unordered condition (such as whether *x* or *y* is a NaN).

The **class** subroutine returns an integer that represents the classification of the floating-point *x* parameter. Since **class** is a reserved key word in C++. The **class** subroutine can not be invoked in a C++ program. The **\_class** subroutine is an interface for C++ program using the **class** subroutine. The interface and the return value for **class** and **\_class** subroutines are identical. The values returned by the **class** subroutine are defined in the **float.h** header file. The return values are the following:

<b>FP_PLUS_NORM</b>	Positive normalized, nonzero <i>x</i>
<b>FP_MINUS_NORM</b>	Negative normalized, nonzero <i>x</i>
<b>FP_PLUS_DENORM</b>	Positive denormalized, nonzero <i>x</i>
<b>FP_MINUS_DENORM</b>	Negative denormalized, nonzero <i>x</i>
<b>FP_PLUS_ZERO</b>	<i>x</i> = +0.0
<b>FP_MINUS_ZERO</b>	<i>x</i> = -0.0
<b>FP_PLUS_INF</b>	<i>x</i> = +INF
<b>FP_MINUS_INF</b>	<i>x</i> = -INF
<b>FP_NANS</b>	<i>x</i> = Signaling Not a Number (NaNS)
<b>FP_NANQ</b>	<i>x</i> = Quiet Not a Number (NaNQ)

Since **class** is a reserved keyword in C++, the **class** subroutine cannot be invoked in a C++ program. The **\_class** subroutine is an interface for the C++ program using the **class** subroutine. The interface and the return values for **class** and **\_class** subroutines are identical.

The **finite** subroutine returns a nonzero value if the *x* parameter is a finite number; that is, if *x* is not +-, INF, NaNQ, or NaNS.

The **isnan** subroutine returns a nonzero value if the *x* parameter is an NaNS or a NaNQ. Otherwise, it returns 0.

The **unordered** subroutine returns a nonzero value if a floating-point comparison between *x* and *y* would be unordered. Otherwise, it returns 0.

**Note:** Compile any routine that uses subroutines from the **libm.a** library with the **-lm** flag. To compile the **class.c** file, for example, enter:

```
cc class.c -lm
```



## Parameters

- x* Specifies some double-precision floating-point value.
- y* Specifies some double-precision floating-point value.

## Error Codes

The **finite**, **isnan**, and **unordered** subroutines neither return errors nor set bits in the floating-point exception status, even if a parameter is an NaNs.

## Related Information

List of Numerical Manipulation Services and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## clock Subroutine

### Purpose

Reports central processing unit (CPU) time used.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <time.h>
clock_t clock (void);
```

### Description

The **clock** subroutine reports the amount of CPU time used. The reported time is the sum of the CPU time of the calling process and its terminated child processes for which it has executed **wait**, **system**, or **pclose** subroutines. To measure the amount of time used by a program, the **clock** subroutine should be called at the beginning of the program, and that return value should be subtracted from the return value of subsequent calls to the **clock** subroutine. To find the time in seconds, divide the value returned by the **clock** subroutine by the value of the macro **CLOCKS\_PER\_SEC**, which is defined in the **time.h** file.

### Return Values

The **clock** subroutine returns the amount of CPU time used.

## Related Information

The **getrusage**, **times** (“getrusage, getrusage64, times, or vtimes Subroutine” on page 428) subroutine, **pclose** (“pclose Subroutine” on page 1002) subroutine, **system** subroutine, **vtimes** (“getrusage, getrusage64, times, or vtimes Subroutine” on page 428) subroutine, **wait**, **waitpid**, **wait3** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## clock\_getcpuclockid Subroutine

### Purpose

Accesses a process CPU-time clock.

## Syntax

```
#include <time.h>
int clock_getcpuclockid(pid_t pid, clockid_t *clock_id);
```

## Description

The **clock\_getcpuclockid** subroutine returns the clock ID of the CPU-time clock of the process specified by *pid*. If the process described by *pid* exists and the calling process has permission, the clock ID of this clock returns in *clock\_id*.

If *pid* is zero, the **clock\_getcpuclockid** subroutine returns the clock ID specified in *clock\_id* of the CPU-time clock of the process making the call.

To obtain the CPU-time clock ID of other processes, the calling process should be root or have the same effective or real user ID as the process that owns the targeted CPU-time clock.

## Parameters

*clock\_id*                Specifies the clock ID of the CPU-time clock.  
*pid*                      Specifies the process ID of the CPU-time clock.

## Return Values

Upon successful completion, the **clock\_getcpuclockid** subroutine returns 0; otherwise, an error code is returned indicating the error.

## Error Codes

**ENOTSUP**                The function is not supported with checkpoint-restart processes.  
**EPERM**                  The requesting process does not have permission to access the CPU-time clock for the process.  
**ESRCH**                  No process can be found corresponding to the process specified by *pid*.

## Related Information

“clock\_getres, clock\_gettime, and clock\_settime Subroutine,” timer\_create subroutine.

---

## clock\_getres, clock\_gettime, and clock\_settime Subroutine

### Purpose

Clock and timer functions.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <time.h>

int clock_getres (clock_id, res)
clockid_t clock_id;
struct timespec *res;

int clock_gettime (clock_id, tp)
clockid_t clock_id;
struct timespec *tp;
```

```
int clock_settime (clock_id, tp)
clockid_t clock_id;
const struct timespec *tp;
```

## Description

The **clock\_getres** subroutine returns the resolution of any clock. Clock resolutions are implementation-defined and cannot be set by a process. If the *res* parameter is not NULL, the resolution of the specified clock is stored in the location pointed to by the *res* parameter. If the *res* parameter is NULL, the clock resolution is not returned. If the *time* parameter of the **clock\_settime** subroutine is not a multiple of the *res* parameter, the value is truncated to a multiple of the *res* parameter.

The **clock\_gettime** subroutine returns the current value, *tp*, for the specified clock, *clock\_id*.

The **clock\_settime** subroutine sets the specified clock, *clock\_id*, to the value specified by the *tp* parameter. Time values that are between two consecutive non-negative integer multiples of the resolution of the specified clock will be truncated down to the smaller multiple of the resolution.

A clock may be system-wide (visible to all processes) or per-process (measuring time that is meaningful only within a process). All implementations support a *clock\_id* of **CLOCK\_REALTIME** as defined in the **time.h** file. This clock represents the Realtime clock for the system. For this clock the values returned by the **clock\_gettime** subroutine and specified by the **clock\_settime** subroutine represent the amount of time (in seconds and nanoseconds) since the epoch.

If the value of the **CLOCK\_REALTIME** clock is set through the **clock\_settime** subroutine, the new value of the clock is used to determine the time of expiration for absolute time services based upon the **CLOCK\_REALTIME** clock. This applies to the time at which armed absolute timers expire. If the absolute time requested at the invocation of such a time service is before the new value of the clock, the time service expires immediately as if the clock had reached the requested time normally.

Setting the value of the **CLOCK\_REALTIME** clock through the **clock\_settime** subroutine has no effect on threads that are blocked waiting for a relative time service based upon this clock, including the **nanosleep** subroutine; nor on the expiration of relative timers based upon this clock. Consequently, these time services expire when the requested relative interval elapses, independently of the new or old value of the clock.

A *clock\_id* of **CLOCK\_MONOTONIC** is defined in the **time.h** file. This clock represents the monotonic clock for the system. For this clock, the value returned by the **clock\_gettime** subroutine represents the amount of time (in seconds and nanoseconds) since an unspecified point in the past. This point does not change after system start time (for example, this clock cannot have backward jumps). The value of the **CLOCK\_MONOTONIC** clock cannot be set through the **clock\_settime** subroutine. This subroutine fails if it is invoked with a *clock\_id* parameter of **CLOCK\_MONOTONIC**.

The calling process should have **SYS\_OPER** authority to set the value of the **CLOCK\_REALTIME** clock.

Process CPU-time clocks are supported by the system. For these clocks, the values returned by **clock\_gettime** and specified by **clock\_settime** represent the amount of execution time of the process associated with the clock. **Clockid\_t** values for CPU-time clocks are obtained by calling **clock\_getcpuclockid**. A special **clockid\_t** value, **CLOCK\_PROCESS\_CPUTIME\_ID**, is defined in the **time.h** file. This value represents the CPU-time clock of the calling process when one of the **clock\_\*** or **timer\_\*** functions is called.

To get or set the value of a CPU-time clock, the calling process must have root permissions or have the same effective or real user ID as the process that owns the targeted CPU-time clock. The same rule applies to a process that tries to get the resolution of a CPU-time clock.

Thread CPU-time clocks are supported by the system. For these clocks, the values returned by **clock\_gettime** and specified by **clock\_settime** represent the amount of execution time of the thread associated with the clock. **Clockid\_t** values for thread CPU-time clocks are obtained by calling the **pthread\_getcpuclockid** subroutine. A special **clockid\_t** value, **CLOCK\_THREAD\_CPUTIME\_ID**, is defined in the **time.h** file. This value represents the thread CPU-time clock of the calling thread when one of the **clock\_\***() or **timer\_\*** functions is called.

To get or set the value of a thread CPU-time clock, the calling thread must be a thread in the same process as the one that owns the targeted thread CPU-time clock. The same rule applies to a thread that tries to get the resolution of a thread CPU-time clock.

## Parameters

<i>clock_id</i>	Specifies the clock.
<i>res</i>	Stores the resolution of the specified clock.
<i>tp</i>	Stores the current value of the specified clock.

## Return Values

If successful, 0 is returned. If unsuccessful, -1 is returned, and **errno** will be set to indicate the error.

## Error Codes

The **clock\_getres**, **clock\_gettime**, and **clock\_settime** subroutines fail if:

<b>EINVAL</b>	The <i>clock_id</i> parameter does not specify a known clock.
<b>ENOTSUP</b>	The function is not supported with checkpoint-restart processes.

The **clock\_settime** subroutine fails if:

<b>EINVAL</b>	The <i>tp</i> parameter to the <b>clock_settime</b> subroutine is outside the range for the given clock ID.
<b>EINVAL</b>	The <i>tp</i> parameter specified a nanosecond value less than zero or greater than or equal to 1000 million.
<b>EINVAL</b>	The value of the <i>clock_id</i> argument is <b>CLOCK_MONOTONIC</b> .

The **clock\_settime** subroutine might fail if:

<b>EPERM</b>	The requesting process does not have the appropriate privilege to set the specified clock.
--------------	--

## Related Information

“clock\_getcpuclockid Subroutine” on page 169, “ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200, “pthread\_getcpuclockid Subroutine” on page 1252, and “nanosleep Subroutine” on page 896.

The **timer\_create** and **timer\_getoverrun** subroutines in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The **time** command in *AIX 5L Version 5.3 Commands Reference, Volume 5*.

---

## clock\_nanosleep Subroutine

### Purpose

Specifies clock for high resolution sleep.

## Syntax

```
#include <time.h>
int clock_nanosleep(clockid_t clock_id, int flags,
    const struct timespec *rqtp, struct timespec *rmtp);
```

## Description

If the **TIMER\_ABSTIME** flag is not set in the *flags* argument, the **clock\_nanosleep** subroutine causes the current thread to be suspended from execution until either the time interval specified by the *rqtp* argument has elapsed, or a signal is delivered to the calling thread and its action is to invoke a signal-catching function, or the process is terminated. The *clock\_id* argument specifies the clock used to measure the time interval.

If the **TIMER\_ABSTIME** flag is set in the *flags* argument, the **clock\_nanosleep** subroutine causes the current thread to be suspended from execution until either the time value of the clock specified by *clock\_id* reaches the absolute time specified by the *rqtp* argument, or a signal is delivered to the calling thread and its action is to invoke a signal-catching function, or the process is terminated. If, at the time of the call, the time value specified by *rqtp* is less than or equal to the time value of the specified clock, then the **clock\_nanosleep** subroutine returns immediately and the calling process shall not be suspended.

The suspension time caused by this function might be longer than requested either because the argument value is rounded up to an integer multiple of the sleep resolution, or because of the scheduling of other activity by the system. Except for the case of being interrupted by a signal, the suspension time for the relative **clock\_nanosleep** subroutine (that is, with the **TIMER\_ABSTIME** flag not set) shall not be less than the time interval specified by the *rqtp* argument, as measured by the corresponding clock. The suspension for the absolute **clock\_nanosleep** subroutine (that is, with the **TIMER\_ABSTIME** flag set) is in effect at least until the value of the corresponding clock reaches the absolute time specified by the *rqtp* argument, except for the case of being interrupted by a signal.

The **clock\_nanosleep** subroutine has no effect on the action or blocking of any signal.

The subroutine fails if the *clock\_id* argument refers to a process or a thread CPU-time clock.

## Parameters

<i>clock_id</i>	Specifies the clock used to measure the time.
<i>flags</i>	Identifies the type of timeout. If <b>TIMER_ABSTIME</b> is set, the time value pointed to by <i>rqtp</i> is an absolute time value; otherwise, it is a time interval.
<i>rmtp</i>	Points to the <b>timespec</b> structure used to return the remaining amount of time in an interval (the requested time minus the time actually slept) if the sleep is interrupted.
<i>rqtp</i>	Points to the <b>timespec</b> structure that contains requested sleep time.

## Return Values

The **clock\_nanosleep** subroutine returns 0 when the requested time has elapsed.

The **clock\_nanosleep** subroutine returns the corresponding error value when it has been interrupted by a signal. For the relative **clock\_nanosleep** subroutine, when the *rmtp* argument is not NULL, the referenced **timespec** structure is updated to contain the amount of time remaining in the interval (the requested time minus the time actually slept). If the *rmtp* argument is NULL, the remaining time is not returned. The absolute **clock\_nanosleep** subroutine has no effect on the structure referenced by the *rmtp* argument.

## Error Codes

**EINTR** The **clock\_nanosleep** subroutine was interrupted by a signal.

<b>EINVAL</b>	The <i>rqtp</i> parameter specified a nanosecond value less than 0 or greater than or equal to 1000 million; or the <b>TIMER_ABSTIME</b> flag was specified in the <b>flags</b> parameter and the <i>rqtp</i> parameter is outside the range for the clock specified by <i>clock_id</i> ; or the <i>clock_id</i> parameter does not specify a known clock, or specifies the CPU-time clock of the calling thread.
<b>ENOTSUP</b>	The <i>clock_id</i> argument specifies a clock for which the <i>clock_nanosleep</i> subroutine is not supported, such as a CPU-time clock.
<b>ENOTSUP</b>	The subroutine is not supported with checkpoint-restarted processes.

## Files

**timer.h**

## Related Information

“clock\_getres, clock\_gettime, and clock\_settime Subroutine” on page 170, “nanosleep Subroutine” on page 896, “pthread\_cond\_wait or pthread\_cond\_timedwait Subroutine” on page 1236, sleep subroutine.

The **timer.h** file.

The Base Definitions volume of IEEE Std 1003.1-2001.

---

## clog, clogf, or clogl Subroutine

### Purpose

Computes the complex natural logarithm.

### Syntax

```
#include <complex.h>
```

```
double complex clog (z)
double complex z;
```

```
float complex clogf (z)
float complex z;
```

```
long double complex clogl (z)
long double complex z;
```

### Description

The **clog**, **clogf**, and **clogl** subroutines compute the complex natural (base *e*) logarithm of *z*, with a branch cut along the negative real axis.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **clog**, **clogf**, and **clogl** subroutines return the complex natural logarithm value, in the range of a strip mathematically unbounded along the real axis and in the interval  $[-i\pi, +i\pi]$  along the imaginary axis.

## Related Information

“cexp, cexpf, or cexpl Subroutine” on page 141

---

## close Subroutine

### Purpose

Closes a file descriptor.

### Syntax

```
#include <unistd.h>

int close (
    FileDescriptor)
int FileDescriptor;
```

### Description

The **close** subroutine closes the file or shared memory object associated with the *FileDescriptor* parameter. If Network File System (NFS) is installed on your system, this file can reside on another node.

All file regions associated with the file specified by the *FileDescriptor* parameter that this process has previously locked with the **lockf** or **fcntl** subroutine are unlocked. This occurs even if the process still has the file open by another file descriptor.

If the *FileDescriptor* parameter resulted from an **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine that specified **O\_DEFER**, and this was the last file descriptor, all changes made to the file since the last **fsync** subroutine are discarded.

If the *FileDescriptor* parameter is associated with a mapped file, it is unmapped. The **shmat** subroutine provides more information about mapped files.

The **close** subroutine attempts to cancel outstanding asynchronous I/O requests on this file descriptor. If the asynchronous I/O requests cannot be canceled, the application is blocked until the requests have completed.

If the *FileDescriptor* parameter is associated with a shared memory object and the shared memory object remains referenced at the last close (that is, a process has it mapped), the entire contents of the memory object persists until the memory object becomes unreferenced. If this is the last close of a shared memory object and the close results in the memory object becoming unreferenced, and the memory object has been unlinked, the memory object is removed. The **shm\_open** subroutine provides more information about shared memory objects.

The **close** subroutine is blocked until all subroutines which use the file descriptor return to **usr** space. For example, when a thread is calling **close** and another thread is calling **select** with the same file descriptor, the **close** subroutine does not return until the **select** call returns.

When all file descriptors associated with a pipe or FIFO special file have been closed, any data remaining in the pipe or FIFO is discarded. If the link count of the file is 0 when all file descriptors associated with the file have been closed, the space occupied by the file is freed, and the file is no longer accessible.

**Note:** If the *FileDescriptor* parameter refers to a device and the **close** subroutine actually results in a device **close**, and the device **close** routine returns an error, the error is returned to the application. However, the *FileDescriptor* parameter is considered closed and it may not be used in any subsequent calls.

All open file descriptors are closed when a process exits. In addition, file descriptors may be closed during the **exec** subroutine if the **close-on-exec** flag has been set for that file descriptor.

## Parameters

*FileDescriptor* Specifies a valid open file descriptor.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to identify the error. If the **close** subroutine is interrupted by a signal that is caught, it returns a value of -1, the **errno** global variable is set to **EINTR** and the state of the *FileDescriptor* parameter is closed.

## Error Codes

The **close** subroutine is unsuccessful if the following is true:

**EBADF** The *FileDescriptor* parameter does not specify a valid open file descriptor.  
**EINTR** Specifies that the **close** subroutine was interrupted by a signal.

The **close** subroutine may also be unsuccessful if the file being closed is NFS-mounted and the server is down under the following conditions:

- The file is on a hard mount.
- The file is locked in any manner.

The **close** subroutine may also be unsuccessful if NFS is installed and the following is true:

**ETIMEDOUT** The connection timed out.

## Related Information

The **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exec Subroutine” on page 236) subroutines, **fcntl** (“fcntl, dup, or dup2 Subroutine” on page 255) subroutine, **ioctl** (“ioctl, ioctlx, ioctl32, or ioctl32x Subroutine” on page 561) subroutine, **lockfx** (“lockfx, lockf, flock, or lockf64 Subroutine” on page 740) subroutine, **open**, **openx**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **pipe** (“pipe Subroutine” on page 1025) subroutine, **socket** subroutine.

The Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## compare\_and\_swap Subroutine

### Purpose

Conditionally updates or returns a single word variable atomically.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/atomic_op.h>
```

```
boolean_t compare_and_swap ( word_addr, old_val_addr, new_val )  
atomic_p word_addr;  
int *old_val_addr;  
int new_val;
```



## Description

The **compare\_and\_swap** subroutine performs an atomic operation which compares the contents of a single word variable with a stored old value. If the values are equal, a new value is stored in the single word variable and **TRUE** is returned; otherwise, the old value is set to the current value of the single word variable and **FALSE** is returned.

The **compare\_and\_swap** subroutine is useful when a word value must be updated only if it has not been changed since it was last read.

**Note:** The word containing the single word variable must be aligned on a full word boundary.

**Note:** If **compare\_and\_swap** is used as a locking primitive, insert an **isync** at the start of any critical sections.

## Parameters

<i>word_addr</i>	Specifies the address of the single word variable.
<i>old_val_addr</i>	Specifies the address of the old value to be checked against (and conditionally updated with) the value of the single word variable.
<i>new_val</i>	Specifies the new value to be conditionally assigned to the single word variable.

## Return Values

<b>TRUE</b>	Indicates that the single word variable was equal to the old value, and has been set to the new value.
<b>FALSE</b>	Indicates that the single word variable was not equal to the old value, and that its current value has been returned in the location where the old value was previously stored.

## Related Information

The **fetch\_and\_add** (“fetch\_and\_add Subroutine” on page 271) subroutine, **fetch\_and\_and** (“fetch\_and\_and or fetch\_and\_or Subroutine” on page 272) subroutine, **fetch\_and\_or** (“fetch\_and\_and or fetch\_and\_or Subroutine” on page 272) subroutine.

---

## compile, step, or advance Subroutine

### Purpose

Compiles and matches regular-expression patterns.

**Note:** Commands use the **regcomp**, **regexec**, **regfree**, and **regerror** subroutines for the functions described in this article.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#define INIT declarations
#define GETC( ) getc_code
#define PEEKC( ) peekc_code
#define UNGETC(c) ungetc_code
#define RETURN(pointer) return_code
#define ERROR(val) error_code
```

```

#include <regexp.h>
#include <NLregexp.h>

char *compile (InString, ExpBuffer, EndBuffer, EndOfFile)
char * ExpBuffer;
char * InString, * EndBuffer;
int EndOfFile;

int step (String, ExpBuffer)
const char * String, *ExpBuffer;

int advance (String, ExpBuffer)
const char *String, *ExpBuffer;

```

## Description

The `/usr/include/regexp.h` file contains subroutines that perform regular-expression pattern matching. Programs that perform regular-expression pattern matching use this source file. Thus, only the `regexp.h` file needs to be changed to maintain regular expression compatibility between programs.

The interface to this file is complex. Programs that include this file define the following six macros before the `#include <regexp.h>` statement. These macros are used by the `compile` subroutine:

<b>INIT</b>	This macro is used for dependent declarations and initializations. It is placed right after the declaration and opening { (left brace) of the <code>compile</code> subroutine. The definition of the <code>INIT</code> buffer must end with a ; (semicolon). <code>INIT</code> is frequently used to set a register variable to point to the beginning of the regular expression so that this register variable can be used in the declarations for the <code>GETC</code> , <code>PEEKC</code> , and <code>UNGETC</code> macros. Otherwise, you can use <code>INIT</code> to declare external variables that <code>GETC</code> , <code>PEEKC</code> , and <code>UNGETC</code> require.
<b>GETC( )</b>	This macro returns the value of the next character in the regular expression pattern. Successive calls to the <code>GETC</code> macro should return successive characters of the pattern.
<b>PEEKC( )</b>	This macro returns the next character in the regular expression. Successive calls to the <code>PEEKC</code> macro should return the same character, which should also be the next character returned by the <code>GETC</code> macro.
<b>UNGETC(c)</b>	This macro causes the parameter <code>c</code> to be returned by the next call to the <code>GETC</code> and <code>PEEKC</code> macros. No more than one character of pushback is ever needed, and this character is guaranteed to be the last character read by the <code>GETC</code> macro. The return value of the <code>UNGETC</code> macro is always ignored.
<b>RETURN(pointer)</b>	This macro is used for normal exit of the <code>compile</code> subroutine. The <code>pointer</code> parameter points to the first character immediately following the compiled regular expression. This is useful for programs that have memory allocation to manage.

## **ERROR(*val*)**

This macro is used for abnormal exit from the **compile** subroutine. It should never contain a **return** statement. The *val* parameter is an error number. The error values and their meanings are:

<b>Error</b>	<b>Meaning</b>
11	Interval end point too large
16	Bad number
25	\ <i>digit</i> out of range
36	Illegal or missing delimiter
41	No remembered search String
42	\ (?\) imbalance
43	Too many \.(
44	More than two numbers given in \{ \}
45	} expected after \.
46	First number exceeds second in \{ \}
49	[ ] imbalance
50	Regular expression overflow
70	Invalid endpoint in range

The **compile** subroutine compiles the regular expression for later use. The *InString* parameter is never used explicitly by the **compile** subroutine, but you can use it in your macros. For example, you can use the **compile** subroutine to pass the string containing the pattern as the *InString* parameter to **compile** and use the **INIT** macro to set a pointer to the beginning of this string. The example in the “Examples” on page 180 section uses this technique. If your macros do not use *InString*, then call **compile** with a value of **((char \*) 0)** for this parameter.

The *ExpBuffer* parameter points to a character array where the compiled regular expression is to be placed. The *EndBuffer* parameter points to the location that immediately follows the character array where the compiled regular expression is to be placed. If the compiled expression cannot fit in (*EndBuffer-ExpBuffer*) bytes, the call **ERROR(50)** is made.

The *EndOfFile* parameter is the character that marks the end of the regular expression. For example, in the **ed** command, this character is usually / (slash).

The **regexp.h** file defines other subroutines that perform actual regular-expression pattern matching. One of these is the **step** subroutine.

The *String* parameter of the **step** subroutine is a pointer to a null-terminated string of characters to be checked for a match.

The *Expbuffer* parameter points to the compiled regular expression, obtained by a call to the **compile** subroutine.

The **step** subroutine returns the value 1 if the given string matches the pattern, and 0 if it does not match. If it matches, then **step** also sets two global character pointers: **loc1**, which points to the first character that matches the pattern, and **loc2**, which points to the character immediately following the last character that matches the pattern. Thus, if the regular expression matches the entire string, **loc1** points to the first character of the *String* parameter and **loc2** points to the null character at the end of the *String* parameter.

The **step** subroutine uses the global variable **circf**, which is set by the **compile** subroutine if the regular expression begins with a ^ (circumflex). If this variable is set, **step** only tries to match the regular expression to the beginning of the string. If you compile more than one regular expression before executing the first one, save the value of **circf** for each compiled expression and set **circf** to that saved value before each call to **step**.

Using the same parameters that were passed to it, the **step** subroutine calls a subroutine named **advance**. The **step** function increments through the *String* parameter and calls the **advance** subroutine until it returns a 1, indicating a match, or until the end of *String* is reached. To constrain the *String* parameter to the beginning of the string in all cases, call the **advance** subroutine directly instead of calling the **step** subroutine.

When the **advance** subroutine encounters an \* (asterisk) or a { \} sequence in the regular expression, it advances its pointer to the string to be matched as far as possible and recursively calls itself, trying to match the rest of the string to the rest of the regular expression. As long as there is no match, the **advance** subroutine backs up along the string until it finds a match or reaches the point in the string that initially matched the \* or { \}. You can stop this backing-up before the initial point in the string is reached. If the **locs** global character is equal to the point in the string sometime during the backing-up process, the **advance** subroutine breaks out of the loop that backs up and returns 0. This is used for global substitutions on the whole line so that expressions such as *s/y\*/g* do not loop forever.

**Note:** In 64-bit mode, these interfaces are not supported: they fail with a return code of 0. In order to use the 64-bit version of this functionality, applications should migrate to the **fnmatch**, **glob**, **regcomp**, and **regex** functions which provide full internationalized regular expression functionality compatible with ISO 9945-1:1996 (IEEE POSIX 1003.1) and with the UNIX98 specification.

## Parameters

<i>InString</i>	Specifies the string containing the pattern to be compiled. The <i>InString</i> parameter is not used explicitly by the <b>compile</b> subroutine, but it may be used in macros.
<i>ExpBuffer</i>	Points to a character array where the compiled regular expression is to be placed.
<i>EndBuffer</i>	Points to the location that immediately follows the character array where the compiled regular expression is to be placed.
<i>EndOfFile</i>	Specifies the character that marks the end of the regular expression.
<i>String</i>	Points to a null-terminated string of characters to be checked for a match.

## Examples

The following is an example of the regular expression macros and calls:

```
#define INIT      register char *sp=instring;
#define GETC()    (*sp++)
#define PEEKC()  (*sp)
#define UNGETC(c) (--sp)
#define RETURN(c) return;
#define ERROR(c) regerr()

#include <regex.h>
. . .
compile (patstr,expbuf, &expbuf[ESIZE], '\0');
. . .
if (step (linebuf, expbuf))
    succeed( );
. . .
```

## Related Information

The **regcmp** or **regex** subroutine, **regcomp** subroutine, **regerror** subroutine, **regex** subroutine, **regfree** subroutine.

List of String Manipulation Services and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## confstr Subroutine

### Purpose

Gets configurable variables.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
size_t confstr (int name, char * buf, size_t len );
```

### Description

The **confstr** subroutine determines the current setting of certain system parameters, limits, or options that are defined by a string value. It is mainly used by applications to find the system default value for the **PATH** environment variable. Its use and purpose are similar to those of the **sysconf** subroutine, but it returns string values rather than numeric values.

If the *Len* parameter is not 0 and the *Name* parameter has a system-defined value, the **confstr** subroutine copies that value into a *Len*-byte buffer pointed to by the *Buf* parameter. If the string returns a value longer than the value specified by the *Len* parameter, including the terminating null byte, then the **confstr** subroutine truncates the string to *Len*-1 bytes and adds a terminating null byte to the result. The application can detect that the string was truncated by comparing the value returned by the **confstr** subroutine with the value specified by the *Len* parameter.

### Parameters

<i>Name</i>	Specifies the system variable setting to be returned. Valid values for the <i>Name</i> parameter are defined in the <b>unistd.h</b> file.
<i>Buf</i>	Points to the buffer into which the <b>confstr</b> subroutine copies the value of the <i>Name</i> parameter.
<i>Len</i>	Specifies the size of the buffer storing the value of the <i>Name</i> parameter.

### Return Values

If the value specified by the *Name* parameter is system-defined, the **confstr** subroutine returns the size of the buffer needed to hold the entire value. If this return value is greater than the value specified by the *Len* parameter, the string returned as the *Buf* parameter is truncated.

If the value of the *Len* parameter is set to 0 and the *Buf* parameter is a null value, the **confstr** subroutine returns the size of the buffer needed to hold the entire system-defined value, but does not copy the string value. If the value of the *Len* parameter is set to 0 but the *Buf* parameter is not a null value, the result is unspecified.

## Error Codes

The **confstr** subroutine will fail if:

**EINVAL**                The value of the name argument is invalid.

## Example

To find out what size buffer is needed to store the string value of the *Name* parameter, enter:

```
confstr(_CS_PATH, NULL, (size_t) 0)
```

The **confstr** subroutine returns the size of the buffer.

## Files

<b>/usr/include/limits.h</b>	Contains system-defined limits.
<b>/usr/include/unistd.h</b>	Contains system-defined environment variables.

## Related Information

The **pathconf** (“pathconf or fpathconf Subroutine” on page 979) subroutine, **sysconf** subroutine.

The **unistd.h** header file.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## conj, conjf, or conjl Subroutine

### Purpose

Computes the complex conjugate.

### Syntax

```
#include <complex.h>
```

```
double complex conj (z)  
double complex z;
```

```
float complex conjf (z)  
float complex z;
```

```
long double complex conjl (z)  
long double complex z;
```

### Description

The **conj**, **conjf**, or **conjl** subroutines compute the complex conjugate of *z*, by reversing the sign of its imaginary part.

### Parameters

*z*                        Specifies the value to be computed.

### Return Values

The **conj**, **conjf**, or **conjl** subroutines return the complex conjugate value.

## Related Information

The “carg, cargf, or cargl Subroutine” on page 131, “cimag, cimagf, or cimagl Subroutine” on page 163, “cproj, cprojf, or cprojl Subroutine” on page 190, “creal, crealf, or creall Subroutine” on page 190.

---

## conv Subroutines

### Purpose

Translates characters.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <ctype.h>
```

```
int toupper ( Character )  
int Character;
```

```
int tolower (Character )  
int Character;
```

```
int _toupper (Character )  
int Character;
```

```
int _tolower (Character )  
int Character;
```

```
int toascii (Character )  
int Character;
```

```
int NCcsc ( Pointer, CharacterPointer )  
NLchar *Pointer;  
char *CharacterPointer;
```

```
int Nctoupper ( Xcharacter )  
int Xcharacter;
```

```
int Nctolower (Xcharacter )  
int Xcharacter;
```

```
int _Nctoupper (Xcharacter )  
int Xcharacter;
```

```
int _Nctolower (Xcharacter )  
int Xcharacter;
```

```
int NctoNLchar (Xcharacter )  
int Xcharacter;
```

```
int NCunescc (CharacterPointer, Pointer )  
char *CharacterPointer;  
NLchar *Pointer;
```

```
int NCflatchr (Xcharacter)
int Xcharacter;
```

## Description

The **toupper** and the **tolower** subroutines have as domain an **int**, which is representable as an unsigned **char** or the value of **EOF**: -1 through 255.

If the parameter of the **toupper** subroutine represents a lowercase letter and there is a corresponding uppercase letter (as defined by **LC\_CTYPE**), the result is the corresponding uppercase letter. If the parameter of the **tolower** subroutine represents an uppercase letter, and there is a corresponding lowercase letter (as defined by **LC\_CTYPE**), the result is the corresponding lowercase letter. All other values in the domain are returned unchanged. If case-conversion information is not defined in the current locale, these subroutines determine character case according to the "C" locale.

The **\_toupper** and **\_tolower** subroutines accomplish the same thing as the **toupper** and **tolower** subroutines, but they have restricted domains. The **\_toupper** routine requires a lowercase letter as its parameter; its result is the corresponding uppercase letter. The **\_tolower** routine requires an uppercase letter as its parameter; its result is the corresponding lowercase letter. Values outside the domain cause undefined results.

The **NCxxxxx** subroutines translate all characters, including extended characters, as code points. The other subroutines translate traditional ASCII characters only. The **NCxxxxx** subroutines are obsolete and should not be used if portability and future compatibility are a concern.

The value of the *Xcharacter* parameter is in the domain of any legal **NLchar** data type. It can also have a special value of -1, which represents the end of file (**EOF**).

If the parameter of the **NCtoupper** subroutine represents a lowercase letter according to the current collating sequence configuration, the result is the corresponding uppercase letter. If the parameter of the **NCtolower** subroutine represents an uppercase letter according to the current collating sequence configuration, the result is the corresponding lowercase letter. All other values in the domain are returned unchanged.

The **\_NCtoupper** and **\_NCtolower** routines are macros that perform the same function as the **NCtoupper** and **NCtolower** subroutines, but have restricted domains and are faster. The **\_NCtoupper** macro requires a lowercase letter as its parameter; its result is the corresponding uppercase letter. The **\_NCtolower** macro requires an uppercase letter as its parameter; its result is the corresponding lowercase letter. Values outside the domain cause undefined results.

The **NCtoNLchar** subroutine yields the value of its parameter with all bits turned off that are not part of an **NLchar** data type.

The **NCesc** subroutine converts the **NLchar** value of the *Pointer* parameter into one or more ASCII bytes stored in the character array pointed to by the *CharacterPointer* parameter. If the **NLchar** data type represents an extended character, it is converted into a printable ASCII escape sequence that uniquely identifies the extended character. **NCesc** returns the number of bytes it wrote. The display symbol table lists the escape sequence for each character.

The opposite conversion is performed by the **NCunes** macro, which translates an ordinary ASCII byte or escape sequence starting at *CharacterPointer* into a single **NLchar** at *Pointer*. **NCunes** returns the number of bytes it read.



The **NCflatchr** subroutine converts its parameter value into the single ASCII byte that most closely resembles the parameter character in appearance. If no ASCII equivalent exists, it converts the parameter value to a ? (question mark).

**Note:** The **setlocale** subroutine may affect the conversion of the decimal point symbol and the thousands separator.

## Parameters

<i>Character</i>	Specifies the character to be converted.
<i>Xcharacter</i>	Specifies an <b>NLchar</b> value to be converted.
<i>CharacterPointer</i>	Specifies a pointer to a single-byte character array.
<i>Pointer</i>	Specifies a pointer to an escape sequence.

## Related Information

The Japanese **conv** (“Japanese conv Subroutines” on page 576) subroutines, **ctype** (“ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines” on page 208) subroutines, **getc**, **fgetc**, **getchar**, or **getw** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **getwc**, **fgetwc**, or **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **setlocale** subroutine.

List of Character Manipulation Services and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## copysign, copysignf, copysignl , copysignd32, copysignd64, and copysignd128 Subroutines

### Purpose

Perform number manipulation.

### Syntax

```
#include <math.h>
```

```
double copysign (x, y)
double x, double y;
```

```
float copysignf (x, y)
float x, float y;
```

```
long double copysignl (x, y)
long double x, long double y;
```

```
_Decimal32 copysignd32(x, y)
_Decimal32 x;
_Decimal32 y;
```

```
_Decimal64 copysignd64(x, y)
_Decimal64 x;
_Decimal64 y;
```

```
_Decimal128 copysignd128(x, y)
_Decimal128 x;
_Decimal128 y;
```

## Description

The **copysign**, **copysignf**, **copysignl**, **copysignd32**, **copysignd64**, and **copysignd128** subroutines produce a value with the magnitude of *x* and the sign of *y*.

## Parameters

*x* Specifies the magnitude.  
*y* Specifies the sign.

## Return Values

Upon successful completion, the **copysign**, **copysignf**, **copysignl**, **copysignd32**, **copysignd64**, and **copysignd128** subroutines return a value with a magnitude of *x* and a sign of *y*.

## Related Information

**signbit** in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## coredump Subroutine

### Purpose

Creates a **core** file without terminating the calling process.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <core.h>
```

```
int coredump( coredumpinfo)  
struct coredumpinfo *coredumpinfo ;
```

### Description

The **coredump** subroutine creates a **core** file of the calling process without terminating the calling process. The created **core** file contains the memory image of the process, and this can be used with the **dbx** command for debugging purposes. In multithreaded processes, only one thread at a time should attempt to call this subroutine. Subsequent calls to **coredump** while a core dump (initiated by another thread) is in progress will fail.

Applications expected to use this facility need to be built with the **-bM:UR binder** flag, otherwise the routine will fail with an error code of **ENOTSUP**.

The **coredumpinfo** structure has the following fields:

Member Type	Member Name	Description
unsigned int	length	Length of the <b>core</b> file name
char *	name	Points to a character string that contains the name of the <b>core</b> file
int	reserved[8]	Reserved fields for future use

## Parameters

*coredumpinfo* Points to the **coredumpinfo** structure

If a NULL pointer is passed as an argument, the default file named **core** in the current directory is used.

## Return Values

Upon successful completion, the **coredump** subroutine returns a value of 0. If the **coredump** subroutine is not successful, a value of -1 is returned and the **errno** global variable is set to indicate the error

## Error Codes

<b>EINVAL</b>	Invalid argument.
<b>EACCES</b>	Search permission is denied on a component of the path prefix, the file exists and the pwrite permission is denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.
<b>EINPROGRESS</b>	A core dump is already in progress.
<b>ENOMEM</b>	Not enough memory.
<b>ENOTSUP</b>	Routine not supported.
<b>EFAULT</b>	Invalid user address.

## Related Information

The adb command, dbx command.

The core file format.

---

## cosf, cosl, or cos Subroutine

### Purpose

Computes the cosine.

### Syntax

```
#include <math.h>
```

```
float cosf (x)  
float x;
```

```
long double cosl (x)  
long double x;
```

```
double cos (x)  
double x;
```

### Description

The **cosf**, **cosl**, and **cos** subroutines compute the cosine of the *x*, parameter (measured in radians).

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **cosf**, **cosl**, and **cos** subroutines return the cosine of *x*.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ , the value 1.0 is returned.

If *x* is  $\pm \text{Inf}$ , a domain error occurs, and a NaN is returned.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

sin, sinl, cos, cosl, tan, or tanl Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## cosh, coshf, or coshl Subroutine

### Purpose

Computes the hyperbolic cosine.

### Syntax

```
#include <math.h>
```

```
float coshf (x)  
float x;
```

```
long double coshl (x)  
long double x;
```

```
double cosh (x)  
double x;
```

### Description

The **coshf**, **coshl**, and **cosh** subroutines compute the hyperbolic cosine of the *x* parameter.

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. On return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **coshf**, **coshl**, and **cosh** subroutines return the hyperbolic cosine of  $x$ .

If the correct value would cause overflow, a range error occurs and the **coshf**, **coshl**, and **cosh** subroutines return the value of the macro **HUGE\_VALF**, **HUGE\_VALL**, and **HUGE\_VAL**, respectively.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$ , the value 1.0 is returned.

If  $x$  is  $\pm \text{Inf}$ ,  $+\text{Inf}$  is returned.

## Related Information

“acosh, acoshf, or acoshl Subroutine” on page 30, “feclearexcept Subroutine” on page 265, “fetetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167

sinh, sinhf, or sinhl Subroutine and tanh, tanhf, or tanhl Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## cpow, cpowf, or cpowl Subroutine

### Purpose

Computes the complex power.

### Syntax

```
#include <complex.h>
```

```
double complex cpow (x, y)
double complex x;
double complex y;
```

```
float complex cpowf (x, y)
float complex x;
float complex y;
```

```
long double complex cpowl (x, y)
long double complex x;
long double complex y;
```

### Description

The **cpow**, **cpowf**, and **cpowl** subroutines compute the complex power function  $x^y$ , with a branch cut for the first parameter along the negative real axis.

### Parameters

$x$  Specifies the base value.  
 $y$  Specifies the power the base value is raised to.

### Return Values

The **cpow**, **cpowf**, and **cpowl** subroutines return the complex power function value.

## Related Information

“cabs, cabsf, or cabsl Subroutine” on page 129 and “csqrt, csqrtf, or csqrtl Subroutine” on page 195

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## cproj, cprojf, or cprojl Subroutine

### Purpose

Computes the complex projection functions.

### Syntax

```
#include <complex.h>
```

```
double complex cproj (z)
double complex z;
```

```
float complex cprojf (z)
float complex z;
```

```
long double complex cprojl (z)
long double complex z;
```

### Description

The **cproj**, **cprojf**, and **cprojl** subroutines compute a projection of  $z$  onto the Riemann sphere:  $z$  projects to  $z$ , except that all complex infinities (even those with one infinite part and one NaN part) project to positive infinity on the real axis. If  $z$  has an infinite part, **cproj**( $z$ ) shall be equivalent to:

```
INFINITY + I * copysign(0.0, cimag(z))
```

### Parameters

$z$  Specifies the value to be projected.

### Return Values

The **cproj**, **cprojf**, and **cprojl** subroutines return the value of the projection onto the Riemann sphere.

## Related Information

“carg, cargf, or cargl Subroutine” on page 131, “cimag, cimagf, or cimagl Subroutine” on page 163, “conj, conjf, or conjl Subroutine” on page 182, and “creal, crealf, or creall Subroutine.”

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## creal, crealf, or creall Subroutine

### Purpose

Computes the real part of a specified value.

### Syntax

```
#include <complex.h>
```

```
double creal (z)
double complex z;
```

```
float crealf (z)
float complex z;
```

```
long double creall (z)
long double complex z;
```

## Description

The **creal**, **crealf**, and **creall** subroutines compute the real part of the value specified by the *z* parameter.

## Parameters

*z* Specifies the real to be computed.

## Return Values

These subroutines return the real part value.

## Related Information

“*carg*, *cargf*, or *cargl* Subroutine” on page 131, “*cimag*, *cimagf*, or *cimagl* Subroutine” on page 163, “*conj*, *conjf*, or *conjl* Subroutine” on page 182, and “*cproj*, *cprojf*, or *cprojl* Subroutine” on page 190

---

## crypt, encrypt, or setkey Subroutine

### Purpose

Encrypts or decrypts data.

### Library

Standard C Library (**libc.a**)

### Syntax

```
char *crypt (PW, Salt)
const char * PW, * Salt;
```

```
void encrypt (Block, EdFlag)
char Block[64];
int EdFlag;
```

```
void setkey (Key)
const char * Key;
```

### Description

The **crypt** and **encrypt** subroutines encrypt or decrypt data. The **crypt** subroutine performs a one-way encryption of a fixed data array with the supplied *PW* parameter. The subroutine uses the *Salt* parameter to vary the encryption algorithm.

The **encrypt** subroutine encrypts or decrypts the data supplied in the *Block* parameter using the key supplied by an earlier call to the **setkey** subroutine. The data in the *Block* parameter on input must be an array of 64 characters. Each character must be an `char 0` or `char 1`.

If you need to statically bind functions from **libc.a** for **crypt** do the following:

1. Create a file and add the following:

```
#!
__setkey
__encrypt
__crypt
```

2. Perform the linking.
3. Add the following to the make file:  
-bI:YourFileName

where *YourFileName* is the name of the file you created in step 1. It should look like the following:  
LDFLAGS=bnoautoimp -bI:/lib/syscalls.exp -bI:YourFileName -lc

These subroutines are provided for compatibility with UNIX system implementations.

## Parameters

<i>Block</i>	Identifies a 64-character array containing the values ( <b>char</b> ) 0 and ( <b>char</b> ) 1. Upon return, this buffer contains the encrypted or decrypted data.
<i>EdFlag</i>	Determines whether the subroutine encrypts or decrypts the data. If this parameter is 0, the data is encrypted. If this is a nonzero value, the data is decrypted. If the <b>/usr/lib/libdes.a</b> file does not exist and the <i>EdFlag</i> parameter is set to nonzero, the <b>encrypt</b> subroutine returns the <b>ENOSYS</b> error code.
<i>Key</i>	Specifies an 64-element array of 0's and 1's cast as a <b>const char</b> data type. The <i>Key</i> parameter is used to encrypt or decrypt data.
<i>PW</i>	Specifies up to an 8-character string to be encrypted.
<i>Salt</i>	Specifies a 2-character string chosen from the following: <b>A-Z</b> Uppercase alpha characters <b>a-z</b> Lowercase alpha characters <b>0-9</b> Numeric characters <b>.</b> Period <b>/</b> Slash The <i>Salt</i> parameter is used to vary the hashing algorithm in one of 4096 different ways.

## Return Values

The **crypt** subroutine returns a pointer to the encrypted password. The static area this pointer indicates may be overwritten by subsequent calls.

## Error Codes

The **encrypt** subroutine returns the following:

<b>ENOSYS</b>	The <b>encrypt</b> subroutine was called with the <i>EdFlag</i> parameter which was set to a nonzero value. Also, the <b>/usr/lib/libdes.a</b> file does not exist.
---------------	---

## Related Information

The **newpass** (“newpass Subroutine” on page 901) subroutine.

The **login** command, **passwd** command, **su** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.



---

## csid Subroutine

### Purpose

Returns the character set ID (charsetID) of a multibyte character.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>

int csid ( String)
const char *String;
```

### Description

The **csid** subroutine returns the charsetID of the multibyte character pointed to by the *String* parameter. No validation of the character is performed. The parameter must point to a value in the character range of the current code set defined in the current locale.

### Parameters

*String*                Specifies the character to be tested.

### Return Values

Successful completion returns an integer value representing the charsetID of the character. This integer can be a number from 0 through *n*, where *n* is the maximum character set defined in the CHARSETID field of the **charmap**. See "Understanding the Character Set Description (charmap) Source File" in *Operating system and device management* for more information.

### Related Information

The **mbstowcs** ("mbstowcs Subroutine" on page 804) subroutine, **wcsid** subroutine.

National Language Support Overview and Understanding the Character Set Description (charmap) Source File in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## csin, csinf, or csinl Subroutine

### Purpose

Computes the complex sine.

### Syntax

```
#include <complex.h>

double complex csin (z)
double complex z;

float complex csinf (z)
```

```
float complex z;
```

```
long double complex csinl (z)  
long double complex z;
```

## Description

The **csin**, **csinf**, and **csinl** subroutines compute the complex sine of the value specified by the *z* parameter.

## Parameters

*z* Specifies the value to be computed.

## Return Values

The **csin**, **csinf**, and **csinl** subroutines return the complex sine value.

## Related Information

“casin, casinf, or casinl Subroutine” on page 131

---

## csinh, csinhf, or csinhl Subroutine

### Purpose

Computes the complex hyperbolic sine.

### Syntax

```
#include <complex.h>
```

```
double complex csinh (z)  
double complex z;
```

```
float complex csinhf (z)  
float complex z;
```

```
long double complex csinhl (z)  
long double complex z;
```

### Description

The **csinh**, **csinhf**, and **csinhl** subroutines compute the complex hyperbolic sine of the value specified by the *z* parameter.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **csinh**, **csinhf**, and **csinhl** subroutines return the complex hyperbolic sine value.

### Related Information

“casinh, casinhf, or casinhl Subroutine” on page 132

---

## csqrt, csqrtf, or csqrtl Subroutine

### Purpose

Computes complex square roots.

### Syntax

```
#include <complex.h>
double complex csqrt (z)
double complex z;

float complex csqrtf (z)
float complex z;

long double complex csqrtl (z)
long double complex z;
```

### Description

The **csqrt**, **csqrtf**, and **csqrtl** subroutines compute the complex square root of the value specified by the *z* parameter, with a branch cut along the negative real axis.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **csqrt**, **csqrtf**, and **csqrtl** subroutines return the complex square root value, in the range of the right half-plane (including the imaginary axis).

### Related Information

“cabs, cabsf, or cabsl Subroutine” on page 129, “cpow, cpowf, or cpowl Subroutine” on page 189

---

## CT\_HOOKx and CT\_GEN macros

### Purpose

Record a trace event into Component Trace, LMT or system trace buffers.

### Syntax

The following set of macros is provided to record a trace entry:

```
#include <sys/ras_trace.h>
CT_HOOK0(ras_block_t cb, int level, int mem_dest, long hkwd);
CT_HOOK1(ras_block_t cb, int level, int mem_dest, long hkwd, long d1);
CT_HOOK2(ras_block_t cb, int level, int mem_dest, long hkwd, long d1, long d2);
CT_HOOK3(ras_block_t cb, int level, int mem_dest, long hkwd, long d1, long d2, long d3);
CT_HOOK4(ras_block_t cb, int level, \
int mem_dest, long hkwd, long d1, long d2, \
long d3, long d4);
CT_HOOK5(ras_block_t cb, int level, int mem_dest, \
long hkwd, long d1, long d2, long d3, \
long d4, long d5);
CT_GEN (ras_block_t cb, int level, long hkwd, long data, long len, void *buf);
```

## Description

The **CT\_HOOKx** macros allow you to record a trace hook. The "x" is the number of data words you want in this trace event.

The **CT\_GEN** macro is used to record a generic trace hook.

All traces are timestamped.

**Restriction:** If the *cb* input parameter has a value of **RAS\_BLOCK\_NULL**, no tracing will be performed.

## Parameters

<b>ras_block_t</b> <i>cb</i>	The <i>cb</i> parameter is the RAS control block that refers to the component that this trace entry belongs to.
<b>int</b> <i>level</i>	<p>The <i>level</i> parameter allows filtering of different trace entries. The higher this level is, the more this trace will be considered as debug or detail information. In other words, this trace entry will appear only if the level of the trace entry is less than or equal to the level of trace chosen for memory or system trace mode.</p> <p>Ten levels of trace are available (<b>CT_LEVEL_0</b> to <b>CT_LEVEL_9</b>, corresponding to value 0 to 9) with four special levels:</p> <ul style="list-style-type: none"><li>• minimal (<b>CT_LVL_MINIMAL</b> (=CT_LEVEL_1))</li><li>• normal (<b>CT_LVL_NORMAL</b> (=CT_LEVEL_3))</li><li>• detail (<b>CT_LVL_DETAIL</b> (=CT_LEVEL_7))</li><li>• default (<b>CT_LVL_DEFAULT</b> (= CT_LVL_MINIMAL (in 5.3 ML 5)))</li></ul> <p>When you are porting an existing driver or subsystem from the existing system trace to component trace, trace existing entries at <b>CT_LVL_DEFAULT</b>.</p>
<b>int</b> <i>mem_dest</i>	<p>For <b>CT_HOOKx</b> macros, the <i>mem_dest</i> parameter indicates the memory destination for this trace entry. It is an ORed value with the following possible settings:</p> <ul style="list-style-type: none"><li>• <b>MT_RARE</b>: the trace entry is saved in the rare buffer of lightweight memory trace if the level condition of the memory trace mode for this control block is satisfied, meaning that the current level of trace for the memory trace mode is greater than or equal to the level of this trace entry.</li><li>• <b>MT_COMMON</b>: the trace entry is saved in the common buffer of the lightweight memory trace if the level condition of the memory trace mode for this control block is satisfied.</li><li>• <b>MT_PRIV</b>: the trace entry is saved in the private memory buffer of the component if the level condition of the memory trace mode for this control block is satisfied.</li><li>• <b>MT_SYSTEM</b>: the trace entry is saved in the existing system trace if the level condition of the <i>system trace mode</i> for this control block is satisfied, if the system trace is running, and if the hook meets any additional criteria specified as part of the system trace. For example, if <b>MT_SYSTEM</b> is not set, the trace entry is not saved in the existing system trace.</li></ul> <p>Only one of the <b>MT_RARE</b>, <b>MT_COMMON</b> and <b>MT_PRIV</b> values should be used, but you can combine ORed with <b>MT_SYSTEM</b>. Otherwise, the trace entry will be duplicated in several memory buffers.</p> <p>The <i>mem_dest</i> parameter is not needed for the <b>CT_GEN</b> macro because lightweight memory trace cannot accommodate generic entries. <b>CT_GEN</b> checks the memory trace and system trace levels to determine whether the generic entry should enter the private memory buffer and system trace buffers respectively.</p>

The *hkwd*, *d1*, *d2*, *d3*, *d4*, *d5*, *len* and *buf* parameters are the same as those used for the existing **TRCHKx** or **TRCGEN** macros. The **TRCHKx** refers to the **TRCHKLnT** macros where *n* is from 0 to 5. For example, **TRCHKL1T** (*hkwd*, *d1*). The **TRCGEN** macros refer to the **TRCGEN** and **TRCGENT** macros.

For the hookword, OR the hookID with a subhookID if needed. For the **CT\_HOOKx** macro, the subhook is ORed into the hookword. For the **CT\_GEN** macro, the subhook is the *d1* parameter.

## Related Information

Trace Facility in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The `trcgenk` and `trcgenkt` kernel services.

The `trchhook`, `trchhook64`, `utrchhook` and `utrchhook64` subroutine.

The `ras_register` and `ras_unregister` exported kernel services.

The `RAS_BLOCK_NULL` Exported Data Structure in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 1*.

---

## CT\_TRCON macro

### Purpose

Return information on whether any trace is active at a certain level for a component.

### Syntax

```
#include <sys/ras_trace.h>
CT_TRCON(cb, level)
```

### Description

The **CT\_TRCON** macro allows you to ascertain whether any type of trace (Component Trace, lightweight memory trace or system trace) will record events for the component specified at the trace detail level specified.

**Note:** If the *cb* input parameter has a value of `RAS_BLOCK_NULL`, the **CT\_TRCON** macro indicates that the trace is off.

### Parameters

<code>ras_block_t</code> <i>cb</i>	The <i>cb</i> parameter is the RAS control block pointer that refers to the component that this trace entry belongs to.
<code>int</code> <i>level</i>	Specifies the trace detail level.

## Related Information

Component Trace Facility in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The “`CT_HOOKx` and `CT_GEN` macros” on page 195.

The `ras_register/ras_unregister` exported kernel services.

The `ras_control` exported kernel services.

The `RAS_BLOCK_NULL` Exported Data Structure in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 1*.

---

## ctan, ctanf, or ctanl Subroutine

### Purpose

Computes complex tangents.

### Syntax

```
#include <complex.h>
```

```
double complex ctan (z)
double complex z;
```

```
float complex ctanf (z)
float complex z;
```

```
long double complex ctanl (z)
long double complex z;
```

### Description

The **ctan**, **ctanf**, and **ctanl** subroutines compute the complex tangent of the value specified by the *z* parameter.

### Parameters

*z* Specifies the value to be computed.

### Return Values

The **ctan**, **ctanf**, and **ctanl** subroutines return the complex tangent value.

### Related Information

“catan, catanf, or catanl Subroutine” on page 132

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## ctanh, ctanhf, or ctanh1 Subroutine

### Purpose

Computes the complex hyperbolic tangent.

### Syntax

```
#include <complex.h>
```

```
double complex ctanh (z)
double complex z;
```

```
float complex ctanhf (z)
float complex z;
```

```
long double complex ctanh1 (z)
long double complex z;
```

### Description

The **ctanh**, **ctanhf**, and **ctanh1** subroutines compute the complex hyperbolic tangent of *z*.

## Parameters

*z* Specifies the value to be computed.

## Return Values

The **ctanh**, **ctanhf**, and **ctanhl** subroutines return the complex hyperbolic tangent value.

## Related Information

“catanh, catanhf, or catanhl Subroutine” on page 133

---

## ctermid Subroutine

### Purpose

Generates the path name of the controlling terminal.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
char *ctermid ( String)
char *String;
```

### Description

The **ctermid** subroutine generates the path name of the controlling terminal for the current process and stores it in a string.

**Note:** File access permissions depend on user access. Access to a file whose path name the **ctermid** subroutine has returned is not guaranteed.

The difference between the **ctermid** and **ttyname** subroutines is that the **ttyname** subroutine must be handed a file descriptor and returns the actual name of the terminal associated with that file descriptor. The **ctermid** subroutine returns a string (the **/dev/tty** file) that refers to the terminal if used as a file name. Thus, the **ttyname** subroutine is useful only if the process already has at least one file open to a terminal.

## Parameters

*String* If the *String* parameter is a null pointer, the string is stored in an internal static area and the address is returned. The next call to the **ctermid** subroutine overwrites the contents of the internal static area.

If the *String* parameter is not a null pointer, it points to a character array of at least `L_ctermid` elements as defined in the **stdio.h** file. The path name is placed in this array and the value of the *String* parameter is returned.

## Related Information

The **isatty** or **ttyname** subroutine.

Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine

### Purpose

Converts the formats of date and time representations.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <time.h>
```

```
char *ctime ( Clock)
const time_t *Clock;
struct tm *localtime (Clock)
const time_t *Clock;
struct tm *gmtime (Clock)
const time_t *Clock;
```

```
time_t mktime( Timeptr)
struct tm *Timeptr;
```

```
double difftime( Time1, Time0)
time_t Time0, Time1;
```

```
char *asctime ( Tm)
const struct tm *Tm;
void tzset ( )
extern long int timezone;
extern int daylight;
extern char *tzname[];
```

### Description

**Attention:** Do not use the **tzset** subroutine when linking with both **libc.a** and **libbsd.a**. The **tzset** subroutine sets the global external variable called **timezone**, which conflicts with the **timezone** subroutine in **libbsd.a**. This name collision may cause unpredictable results.

**Attention:** Do not use the **ctime**, **localtime**, **gmtime**, or **asctime** subroutine in a multithreaded environment. See the multithread alternatives in the **ctime\_r** (“**ctime\_r**, **localtime\_r**, **gmtime\_r**, or **asctime\_r** Subroutine” on page 206), **localtime\_r**, **gmtime\_r**, or **asctime\_r** subroutine article.

The **ctime** subroutine converts a time value pointed to by the *Clock* parameter, which represents the time in seconds since 00:00:00 Coordinated Universal Time (UTC), January 1, 1970, into a 26-character string in the following form:

```
Sun Sept 16 01:03:52 1973\n\0
```

The width of each field is always the same as shown here.

The **ctime** subroutine adjusts for the time zone and daylight saving time, if it is in effect.

The **localtime** subroutine converts the long integer pointed to by the *Clock* parameter, which contains the time in seconds since 00:00:00 UTC, 1 January 1970, into a **tm** structure. The **localtime** subroutine adjusts for the time zone and for daylight-saving time, if it is in effect. Use the time-zone information as though **localtime** called **tzset**.



The **gmtime** subroutine converts the long integer pointed to by the *Clock* parameter into a **tm** structure containing the Coordinated Universal Time (UTC), which is the time standard the operating system uses.

**Note:** UTC is the international time standard intended to replace GMT.

The **tm** structure is defined in the **time.h** file, and it contains the following members:

```
int tm_sec;    /* Seconds (0 - 59) */
int tm_min;    /* Minutes (0 - 59) */
int tm_hour;   /* Hours (0 - 23) */
int tm_mday;   /* Day of month (1 - 31) */
int tm_mon;    /* Month of year (0 - 11) */
int tm_year;   /* Year - 1900 */
int tm_wday;   /* Day of week (Sunday = 0) */
int tm_yday;   /* Day of year (0 - 365) */
int tm_isdst;  /* Nonzero = Daylight saving time */
```

The **mktime** subroutine is the reverse function of the **localtime** subroutine. The **mktime** subroutine converts the **tm** structure into the time in seconds since 00:00:00 UTC, 1 January 1970. The **tm\_wday** and **tm\_yday** fields are ignored, and the other components of the **tm** structure are not restricted to the ranges specified in the **/usr/include/time.h** file. The value of the **tm\_isdst** field determines the following actions of the **mktime** subroutine:

- 0       Initially presumes that Daylight Savings Time (DST) is not in effect.
- >0      Initially presumes that DST is in effect.
- 1      Actively determines whether DST is in effect from the specified time and the local time zone. Local time zone information is set by the **tzset** subroutine.

Upon successful completion, the **mktime** subroutine sets the values of the **tm\_wday** and **tm\_yday** fields appropriately. Other fields are set to represent the specified time since January 1, 1970. However, the values are forced to the ranges specified in the **/usr/include/time.h** file. The final value of the **tm\_mday** field is not set until the values of the **tm\_mon** and **tm\_year** fields are determined.

**Note:** The **mktime** subroutine cannot convert time values before 00:00:00 UTC, January 1, 1970 and after 03:14:07 UTC, January 19, 2038.

The **difftime** subroutine computes the difference between two calendar times: the *Time1* and *-Time0* parameters.

The **asctime** subroutine converts a **tm** structure to a 26-character string of the same format as **ctime**.

If the **TZ** environment variable is defined, then its value overrides the default time zone, which is the U.S. Eastern time zone. The **environment** facility contains the format of the time zone information specified by **TZ**. **TZ** is usually set when the system is started with the value that is defined in either the **/etc/environment** or **/etc/profile** files. However, it can also be set by the user as a regular environment variable for performing alternate time zone conversions.

The **tzset** subroutine sets the **timezone**, **daylight**, and **tzname** external variables to reflect the setting of **TZ**. The **tzset** subroutine is called by **ctime** and **localtime**, and it can also be called explicitly by an application program.

The **timezone** external variable contains the difference, in seconds, between UTC and local standard time. For example, the value of **timezone** is 5 \* 60 \* 60 for U.S. Eastern Standard Time.

The **daylight** external variable is nonzero when a daylight-saving time conversion should be applied. By default, this conversion follows the standard U.S. conventions; other conventions can be specified. The default conversion algorithm adjusts for the peculiarities of U.S. daylight saving time in 1974 and 1975.

The **tzname** external variable contains the name of the standard time zone (**tzname[0]**) and of the time zone when Daylight Savings Time is in effect (**tzname[1]**). For example:

```
char *tzname[2] = {"EST", "EDT"};
```

The **time.h** file contains declarations of all these subroutines and externals and the **tm** structure.

## Parameters

<i>Clock</i>	Specifies the pointer to the time value in seconds.
<i>Timeptr</i>	Specifies the pointer to a <b>tm</b> structure.
<i>Time1</i>	Specifies the pointer to a <b>time_t</b> structure.
<i>Time0</i>	Specifies the pointer to a <b>time_t</b> structure.
<i>Tm</i>	Specifies the pointer to a <b>tm</b> structure.

## Return Values

**Attention:** The return values point to static data that is overwritten by each call.

The **tzset** subroutine returns no value.

The **mktime** subroutine returns the specified time in seconds encoded as a value of type **time\_t**. If the time cannot be represented, the function returns the value **(time\_t)-1**.

The **localtime** and **gmtime** subroutines return a pointer to the **struct tm**.

The **ctime** and **asctime** subroutines return a pointer to a 26-character string.

The **difftime** subroutine returns the difference expressed in seconds as a value of type **double**.

## Related Information

The **getenv** (“getenv Subroutine” on page 365) subroutine, **gettimer** (“gettimer, settimer, restimer, stime, or time Subroutine” on page 446) subroutine, **strftime** subroutine.

Time data manipulation services in *Operating system and device management*.

National Language Support Overview for Programming in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ctime64, localtime64, gmtime64, mktime64, difftime64, or asctime64 Subroutine

### Purpose

Converts the formats of date and time representations.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <time.h>

char *ctime64 (Clock)
const time64_t *Clock;

struct tm *localtime64 (Clock)
const time64_t *Clock;

struct tm *gmtime64 (Clock)
const time64_t *Clock;

time64_t mktime64(Timeptr)
struct tm *Timeptr;

double difftime64(Time1, Time0)
time64_t Time0, Time1;

char *asctime64 (Tm)
const struct tm *Tm;
```

## Description

**Attention:** Do not use the **ctime**, **localtime**, **gmtime**, or **asctime** subroutine in a multithreaded environment. See “**ctime64\_r**, **localtime64\_r**, **gmtime64\_r**, or **asctime64\_r** Subroutine” on page 204 for multithread alternatives.

The **ctime64** subroutine converts a time value pointed to by the *Clock* parameter, which represents the time in seconds since 00:00:00 Coordinated Universal Time (UTC), January 1, 1970, into a 26-character string in the following form:

```
Sun Sept 16 01:03:52 1973\n\0
```

The width of each field is always the same as shown here.

The **ctime64** subroutine adjusts for the time zone and daylight saving time, if it is in effect.

The **localtime64** subroutine converts the 64 bit long pointed to by the *Clock* parameter, which contains the time in seconds since 00:00:00 UTC, 1 January 1970, into a *tm* structure. The **localtime64** subroutine adjusts for the time zone and for daylight saving time, if it is in effect. Use the time-zone information as though **localtime64** called **tzset**.

The **gmtime64** subroutine converts the 64 bit long pointed to by the *Clock* parameter into a *tm* structure containing the Coordinated Universal Time (UTC), which is the time standard that the operating system uses.

**Note:** UTC is the international time standard intended to replace GMT.

The **mktime64** subroutine is the reverse function of the **localtime64** subroutine. The **mktime64** subroutine converts the *tm* structure into the time in seconds since 00:00:00 UTC, 1 January 1970. The **tm\_wday** and **tm\_yday** fields are ignored, and the other components of the *tm* structure are not restricted to the ranges specified in the **/usr/include/time.h** file. The value of the **tm\_isdst** field determines the following actions of the **mktime64** subroutine:

0	Initially presumes that Daylight Savings Time (DST) is not in effect.
>0	Initially presumes that DST is in effect.
-1	Actively determines whether DST is in effect from the specified time and the local time zone. Local time zone information is set by the <b>tzset</b> subroutine.

Upon successful completion, the **mktime64** subroutine sets the values of the **tm\_wday** and **tm\_yday** fields appropriately. Other fields are set to represent the specified time since January 1, 1970. However, the values are forced to the ranges specified in the **/usr/include/time.h** file. The final value of the **tm\_mday** field is not set until the values of the **tm\_mon** and **tm\_year** fields are determined.

**Note:** The **mktime64** subroutine cannot convert time values before 00:00:00 UTC, January 1, 1970 and after 23:59:59 UTC, December 31, 9999.

**Note:** The **difftime64** subroutine computes the difference between two calendar times: the *Time1* and *Time0* parameters.

**Note:** The **asctime64** subroutine converts a *tm* structure to a 26-character string of the same format as **ctime64**.

## Parameters

<i>Clock</i>	Specifies the pointer to the time value in seconds.
<i>Timeptr</i>	Specifies the pointer to a <b>tm</b> structure.
<i>Time1</i>	Specifies the pointer to a <b>time64_t</b> structure.
<i>Time0</i>	Specifies the pointer to a <b>time64_t</b> structure.
<i>Tm</i>	Specifies the pointer to a <b>tm</b> structure.

## Return Values

**Attention:** The return values point to static data that is overwritten by each call.

The **mktime64** subroutine returns the specified time in seconds encoded as a value of type **time64\_t**. If the time cannot be represented, the function returns the value **(time64\_t)-1**.

The **localtime64** and **gmtime64** subroutines return a pointer to the **tm** struct .

The **ctime64** and **asctime64** subroutines return a pointer to a 26-character string.

The **difftime64** subroutine returns the difference expressed in seconds as a value of type long double.

## Related Information

“**ctime64\_r**, **localtime64\_r**, **gmtime64\_r**, or **asctime64\_r** Subroutine,” “**getenv** Subroutine” on page 365, “**gettimer**, **settimer**, **restimer**, **stime**, or **time** Subroutine” on page 446, **strftime** subroutine.

Time data manipulation services in *Operating system and device management*.

National Language Support Overview for Programming in AIX 5L™ Version 5.3 National Language Support Guide and Reference.

Subroutines, Example Programs, and Libraries in AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs.

---

## **ctime64\_r**, **localtime64\_r**, **gmtime64\_r**, or **asctime64\_r** Subroutine

### Purpose

Converts the formats of date and time representations.

## Library

Thread-Safe C Library (**libc\_r.a**)

## Syntax

```
#include <time.h>
```

```
char *ctime64_r(Timer, BufferPointer)
const time64_t * Timer;
char * BufferPointer;
```

```
struct tm *localtime64_r(Timer, CurrentTime)
const time64_t * Timer;
struct tm * CurrentTime;
```

```
struct tm *gmtime64_r (Timer, XTime)
const time64_t * Timer;
struct tm * XTime;
```

```
char *asctime64_r (TimePointer, BufferPointer)
const struct tm * TimePointer;
char * BufferPointer;
```

## Description

The **ctime64\_r** subroutine converts a time value pointed to by the *Timer* parameter, which represents the time in seconds since 00:00:00 Coordinated Universal Time (UTC), January 1, 1970, into the character array pointed to by the *BufferPointer* parameter. The character array should have a length of at least 26 characters so the converted time value fits without truncation. The converted time value string takes the form of the following example:

```
Sun Sept 16 01:03:52 1973\n\0
```

The width of each field is always the same as shown here. Thus, *ctime* will only return dates up to December 31, 9999.

The **ctime64\_r** subroutine adjusts for the time zone and daylight saving time, if it is in effect.

The **localtime64\_r** subroutine converts the **time64\_t** structure pointed to by the *Timer* parameter, which contains the time in seconds since 00:00:00 UTC, January 1, 1970, into the **tm** structure pointed to by the *CurrentTime* parameter. The **localtime64\_r** subroutine adjusts for the time zone and for daylight saving time, if it is in effect.

The **gmtime64\_r** subroutine converts the **time64\_t** structure pointed to by the *Timer* parameter into the **tm** structure pointed to by the *XTime* parameter.

The **tm** structure is defined in the **time.h** header file. The **time.h** file contains declarations of these subroutines, externals, and the **tm** structure.

The **asctime64\_r** subroutine converts the **tm** structure pointed to by the *TimePointer* parameter into a 26-character string in the same format as the **ctime64\_r** subroutine. The results are placed into the character array, *BufferPointer*. The *BufferPointer* parameter points to the resulting character array, which takes the form of the following example:

```
Sun Sept 16 01:03:52 1973\n\0
```

Programs using this subroutine must link to the **libpthreads.a** library.

## Parameters

<i>Timer</i>	Points to a <b>time64_t</b> structure, which contains the number of seconds since 00:00:00 UTC, January 1, 1970.
<i>BufferPointer</i>	Points to a character array at least 26 characters long.
<i>CurrentTime</i>	Points to a <b>tm</b> structure. The result of the <b>localtime64_r</b> subroutine is placed here.
<i>XTime</i>	Points to a <b>tm</b> structure used for the results of the <b>gmtime64_r</b> subroutine.
<i>TimePointer</i>	Points to a <b>tm</b> structure used as input to the <b>asctime64_r</b> subroutine.

## Return Values

The **localtime64\_r** and **gmtime64\_r** subroutines return a pointer to the **tm** structure. The **asctime64\_r** returns NULL if either *TimePointer* or *BufferPointer* is NULL.

The **ctime64\_r** and **asctime64\_r** subroutines return a pointer to a 26-character string. The **ctime64\_r** subroutine returns NULL if the *BufferPointer* is NULL.

The **difftime64** subroutine returns the difference expressed in seconds as a value of type long double.

## Files

**/usr/include/time.h** Defines time macros, data types, and structures.

## Related Information

“ctime64, localtime64, gmtime64, mktime64, difftime64, or asctime64 Subroutine” on page 202

Subroutines, Example Programs, and Libraries in AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs.

---

## ctime\_r, localtime\_r, gmtime\_r, or asctime\_r Subroutine

### Purpose

Converts the formats of date and time representations.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
#include <time.h>
```

```
char *ctime_r(Timer, BufferPointer)
const time_t * Timer;
char * BufferPointer;
```

```
struct tm *localtime_r(Timer, CurrentTime)
const time_t * Timer;
struct tm * CurrentTime;
```

```
struct tm *gmtime_r(Timer, XTime)
const time_t * Timer;
struct tm * XTime;
```

```

char *asctime_r(TimePointer, BufferPointer)
const struct tm * TimePointer;
char * BufferPointer;

```

## Description

The **ctime\_r** subroutine converts a time value pointed to by the *Timer* parameter, which represents the time in seconds since 00:00:00 Coordinated Universal Time (UTC), January 1, 1970, into the character array pointed to by the *BufferPointer* parameter. The character array should have a length of at least 26 characters so the converted time value fits without truncation. The converted time value string takes the form of the following example:

```
Sun Sep 16 01:03:52 1973\n\0
```

The width of each field is always the same as shown here.

The **ctime\_r** subroutine adjusts for the time zone and daylight saving time, if it is in effect.

The **localtime\_r** subroutine converts the **time\_t** structure pointed to by the *Timer* parameter, which contains the time in seconds since 00:00:00 UTC, January 1, 1970, into the **tm** structure pointed to by the *CurrentTime* parameter. The **localtime\_r** subroutine adjusts for the time zone and for daylight saving time, if it is in effect.

The **gmtime\_r** subroutine converts the **time\_t** structure pointed to by the *Timer* parameter into the **tm** structure pointed to by the *XTime* parameter.

The **tm** structure is defined in the **time.h** header file. The **time.h** file contains declarations of these subroutines, externals, and the **tm** structure.

The **asctime\_r** subroutine converts the **tm** structure pointed to by the *TimePointer* parameter into a 26-character string in the same format as the **ctime\_r** subroutine. The results are placed into the character array, *BufferPointer*. The *BufferPointer* parameter points to the resulting character array, which takes the form of the following example:

```
Sun Sep 16 01:03:52 1973\n\0
```

Programs using this subroutine must link to the **libpthreads.a** library.

## Parameters

<i>Timer</i>	Points to a <b>time_t</b> structure, which contains the number of seconds since 00:00:00 UTC, January 1, 1970.
<i>BufferPointer</i>	Points to a character array at least 26 characters long.
<i>CurrentTime</i>	Points to a <b>tm</b> structure. The result of the <b>localtime_r</b> subroutine is placed here.
<i>XTime</i>	Points to a <b>tm</b> structure used for the results of the <b>gmtime_r</b> subroutine.
<i>TimePointer</i>	Points to a <b>tm</b> structure used as input to the <b>asctime_r</b> subroutine.

## Return Values

The **localtime\_r** and **gmtime\_r** subroutines return a pointer to the **tm** structure. The **asctime\_r** returns NULL if either *TimePointer* or *BufferPointer* are NULL.

The **ctime\_r** and **asctime\_r** subroutines return a pointer to a 26-character string. The **ctime\_r** subroutine returns NULL if the *BufferPointer* is NULL.

## Files

`/usr/include/time.h`

Defines time macros, data types, and structures.

## Related Information

The **ctime**, **localtime**, **gmtime**, **mktime**, **difftime**, **asctime**, or **tzset** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine.

Subroutines, Example Programs, and Libraries and List of Multi-threaded Programming Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*

---

## **ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines**

### Purpose

Classifies characters.

### Library

Standard Character Library (**libc.a**)

### Syntax

```
#include <ctype.h>
```

```
int isalpha ( Character)
```

```
int Character;
```

```
int isupper (Character)
```

```
int Character;
```

```
int islower (Character)
```

```
int Character;
```

```
int isdigit (Character)
```

```
int Character;
```

```
int isxdigit (Character)
```

```
int Character;
```

```
int isalnum (Character)
```

```
int Character;
```

```
int isspace (Character)
```

```
int Character;
```

```
int ispunct (Character)
```

```
int Character;
```

```
int isprint (Character)
```

```
int Character;
```

```
int isgraph (Character)
```

```
int Character;
```

```
int iscntrl (Character)
```

```
int Character;
```

```
int isascii (Character)
```

```
int Character;
```



## Description

The **ctype** subroutines classify character-coded integer values specified in a table. Each of these subroutines returns a nonzero value for True and 0 for False.

**Note:** The **ctype** subroutines should only be used on character data that can be represented by a single byte value ( 0 through 255 ). Attempting to use the **ctype** subroutines on multi-byte locale data may give inconsistent results. Wide character classification routines (such as **iswprint**, **iswlower**, etc.) should be used with dealing with multi-byte character data.

## Locale Dependent Character Tests

The following subroutines return nonzero (True) based upon the character class definitions for the current locale.

<b>isalnum</b>	Returns nonzero for any character for which the <b>isalpha</b> or <b>isdigit</b> subroutine would return nonzero. The <b>isalnum</b> subroutine tests whether the character is of the <b>alpha</b> or <b>digit</b> class.
<b>isalpha</b>	Returns nonzero for any character for which the <b>isupper</b> or <b>islower</b> subroutines would return nonzero. The <b>isalpha</b> subroutine also returns nonzero for any character defined as an alphabetic character in the current locale, or for a character for which <i>none</i> of the <b>isctrl</b> , <b>isdigit</b> , <b>ispunct</b> , or <b>isspace</b> subroutines would return nonzero. The <b>isalpha</b> subroutine tests whether the character is of the <b>alpha</b> class.
<b>isupper</b>	Returns nonzero for any uppercase letter [A through Z]. The <b>isupper</b> subroutine also returns nonzero for any character defined to be uppercase in the current locale. The <b>isupper</b> subroutine tests whether the character is of the <b>upper</b> class.
<b>islower</b>	Returns nonzero for any lowercase letter [a through z]. The <b>islower</b> subroutine also returns nonzero for any character defined to be lowercase in the current locale. The <b>islower</b> subroutine tests whether the character is of the <b>lower</b> class.
<b>isspace</b>	Returns nonzero for any white-space character (space, form feed, new line, carriage return, horizontal tab or vertical tab). The <b>isspace</b> subroutine tests whether the character is of the <b>space</b> class.
<b>ispunct</b>	Returns nonzero for any character for which the <b>isprint</b> subroutine returns nonzero, except the space character and any character for which the <b>isalnum</b> subroutine would return nonzero. The <b>ispunct</b> subroutine also returns nonzero for any locale-defined character specified as a punctuation character. The <b>ispunct</b> subroutine tests whether the character is of the <b>punct</b> class.
<b>isprint</b>	Returns nonzero for any printing character. Returns nonzero for any locale-defined character that is designated a printing character. This routine tests whether the character is of the <b>print</b> class.
<b>isgraph</b>	Returns nonzero for any character for which the <b>isprint</b> character returns nonzero, except the space character. The <b>isgraph</b> subroutine tests whether the character is of the <b>graph</b> class.
<b>isctrl</b>	Returns nonzero for any character for which the <b>isprint</b> subroutine returns a value of False (0) and any character that is designated a control character in the current locale. For the C locale, control characters are the ASCII delete character (0177 or 0x7F), or an ordinary control character (less than 040 or 0x20). The <b>isctrl</b> subroutine tests whether the character is of the <b>ctrl</b> class.

## Locale Independent Character Tests

The following subroutines return nonzero for the same characters, regardless of the locale:

<b>isdigit</b>	<i>Character</i> is a digit in the range [0 through 9].
<b>isxdigit</b>	<i>Character</i> is a hexadecimal digit in the range [0 through 9], [A through F], or [a through f].
<b>isascii</b>	<i>Character</i> is an ASCII character whose value is in the range 0 through 0177 (0 through 0x7F), inclusive.

## Parameter

*Character* Indicates the character to be tested (integer value).

## Return Codes

The **ctype** subroutines return nonzero (True) if the character specified by the *Character* parameter is a member of the selected character class; otherwise, a 0 (False) is returned.

## Related Information

The **setlocale** subroutine.

List of Character Manipulation Services and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## cuserid Subroutine

### Purpose

Gets the alphanumeric user name associated with the current process.

### Library

Standard C Library (**libc.a**)

Use the **libc\_r.a** library to access the thread-safe version of this subroutine.

### Syntax

```
#include <stdio.h>
```

```
char *cuserid ( Name )  
char *Name;
```

### Description

The **cuserid** subroutine gets the alphanumeric user name associated with the current process. This subroutine generates a character string representing the name of a process's owner.

**Note:** The **cuserid** subroutine duplicates functionality available with the **getpwuid** and **getuid** subroutines. Present applications should use the **getpwuid** and **getuid** subroutines.

If the *Name* parameter is a null pointer, then a character string of size `L_cuserid` is dynamically allocated with **malloc**, and the character string representing the name of the process owner is stored in this area. The **cuserid** subroutine then returns the address of this area. Multithreaded application programs should use this functionality to obtain thread specific data, and then continue to use this pointer in subsequent calls to the **cuserid** subroutine. In any case, the application program must deallocate any dynamically allocated space with the **free** subroutine when the data is no longer needed.

If the *Name* parameter is not a null pointer, the character string is stored into the array pointed to by the *Name* parameter. This array must contain at least the number of characters specified by the constant `L_cuserid`. This constant is defined in the **stdio.h** file.

If the user name cannot be found, the **cuserid** subroutine returns a null pointer; if the *Name* parameter is not a null pointer, a null character (`'\0'`) is stored in *Name* [0].

## Parameter

*Name*            Points to a character string representing a user name.

## Related Information

The **endpwent** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422) subroutine, **getlogin** (“getlogin Subroutine” on page 393), **getpwent** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422), **getpwnam** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422), **getpwuid** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422), or **putpwent** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422) subroutine.

Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## defssys Subroutine

### Purpose

Initializes the **SRCsubsys** structure with default values.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <sys/srcobj.h>
#include <spc.h>
```

```
void defssys( SRCSubsystem)
struct SRCsubsys *SRCSubsystem;
```

### Description

The **defssys** subroutine initializes the **SRCsubsys** structure of the **/usr/include/sys/srcobj.h** file with the following default values:

Field	Value
display	SRCYES
multi	SRCNO
contact	SRC SOCKET
waittime	TIMELIMIT
priority	20
action	ONCE
standerr	<b>/dev/console</b>
standin	<b>/dev/console</b>
standout	<b>/dev/console</b>

All other numeric fields are set to 0, and all other alphabetic fields are set to an empty string.

This function must be called to initialize the **SRCsubsys** structure before an application program uses this structure to add records to the subsystem object class.

## Parameters

*SRCSubsystem* Points to the **SRCSubsys** structure.

## Related Information

The **addssys** (“addssys Subroutine” on page 33) subroutine.

Defining Your Subsystem to the SRC, List of SRC Subroutines, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## delssys Subroutine

### Purpose

Removes the subsystem objects associated with the *SubsystemName* parameter.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <sys/srcobj.h>
#include <spc.h>
```

```
int delssys ( SubsystemName )
char *SubsystemName;
```

### Description

The **delssys** subroutine removes the subsystem objects associated with the specified subsystem. This removes all objects associated with that subsystem from the following object classes:

- Subsystem
- Subserver Type
- Notify

The program running with this subroutine must be running with the group **system**.

### Parameter

*SubsystemName* Specifies the name of the subsystem.

### Return Values

Upon successful completion, the **delssys** subroutine returns a positive value. If no record is found, a value of 0 is returned. Otherwise, -1 is returned and the **odmerrno** variable is set to indicate the error. See “Appendix B. ODM Error Codes (Appendix B, “ODM Error Codes,” on page 1347)” for a description of possible **odmerrno** values.

### Security

Privilege Control:

**SET\_PROC\_AUDIT** kernel privilege

Files Accessed:

Mode	File
644	/etc/objrepos/SRCsubsys
644	/etc/objrepos/SRCsubsvr
644	/etc/objrepos/SRCnotify

Auditing Events:

Event	Information
SRC_Delssys	Lists in an audit log the name of the subsystem being removed.

## Files

/etc/objrepos/SRCsubsys	SRC Subsystem Configuration object class.
/etc/objrepos/SRCsubsvr	SRC Subsystem Configuration object class.
/etc/objrepos/SRCnotify	SRC Notify Method object class.
/dev/SRC	Specifies the <b>AF_UNIX</b> socket file.
/dev/.SRC-unix	Specifies the location for temporary socket files.
/usr/include/sys/srcobj.h	Defines object structures used by the SRC.
/usr/include/spc.h	Defines external interfaces provided by the SRC subroutines.

## Related Information

The **addssys** (“addssys Subroutine” on page 33) subroutine, **chssys** (“chssys Subroutine” on page 162) subroutine.

The **chssys** command, **mkssys** command, **rmssys** command.

List of SRC Subroutines and System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## dirname Subroutine

### Purpose

Report the parent directory name of a file path name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <libgen.h>
```

```
char *dirname (path)  
char *path
```

### Description

Given a pointer to a character string that contains a file system path name, the **dirname** subroutine returns a pointer to a string that is the parent directory of that file. Trailing “/” characters in the path are not counted as part of the path.

If *path* is a null pointer or points to an empty string, a pointer to a static constant “.” is returned.

The **dirname** and **basename** subroutines together yield a complete path name. **dirname** (*path*) is the directory where **basename** (*path*) is found.

## Parameters

*path*                    Character string containing a file system path name.

## Return Values

The **dirname** subroutine returns a pointer to a string that is the parent directory of *path*. If *path* or *\*path* is a null pointer or points to an empty string, a pointer to a string "." is returned. The **dirname** subroutine may modify the string pointed to by *path* and may return a pointer to static storage that may then be overwritten by sequent calls to the **dirname** subroutine.

## Examples

A simple file name and the strings "." and ".." all have "." as their return value.

Input string	Output string
/usr/lib	/usr
/usr/	/
usr	.
/	/
.	.
..	.

The following code reads a path name, changes directory to the appropriate directory, and opens the file.

```
char path [MAXPATHEN], *pathcopy;
int fd;
fgets (path, MAXPATHEN, stdin);
pathcopy = strdup (path);
chdir (dirname (pathcopy) );
fd = open (basename (path), O_RDONLY);
```

## Related Information

The **basename** ("basename Subroutine" on page 117) or **chdir** ("chdir Subroutine" on page 147) subroutine.

---

## disclaim and disclaim64 Subroutines

### Purpose

Disclaim the content of a memory address range.

### Syntax

```
#include <sys/shm.h>
int disclaim ( Address, Length, Flag)
char *Address;
unsigned int Length, Flag;
```

```
int disclaim64( Address, Length, Flag)
void *Address;
size_t Length;
unsigned long Flag;
```

## Description

The **disclaim** and **disclaim64** subroutines mark an area of memory having content that is no longer needed. The system then stops paging the memory area. These subroutines cannot be used on memory that is mapped to a file by the **shmat** subroutine.

## Parameters

<i>Address</i>	Points to the beginning of the memory area.
<i>Length</i>	Specifies the length of the memory area in bytes.
<i>Flag</i>	Must be the <b>DISCLAIM_ZEROMEM</b> value, which indicates that each memory location in the address range should be set to zero.

## Return Values

When successful, the **disclaim** and **disclaim64** subroutines return a value of 0.

## Error Codes

If the **disclaim** and **disclaim64** subroutines are not successful, they return a value of -1 and set the **errno** global variable to indicate the error. The **disclaim** and **disclaim64** subroutines are not successful if one or more of the following are true:

<b>EFAULT</b>	The calling process does not have write access to the area of memory that begins at the <i>Address</i> parameter and extends for the number of bytes specified by the <i>Length</i> parameter.
<b>EINVAL</b>	The value of the <i>Flag</i> parameter is not valid.
<b>EINVAL</b>	The memory area is mapped to a file.

## Related Information

The **shm.h** file in *AIX 5L Version 5.3 Files Reference*.

The **shmat** and **shmctl** subroutines in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The **List of Memory Manipulation Services** in the *AIX 5L Version 5.3 General Programming Concepts*.

The **System Calls Available to Kernel Extensions** in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

The **Allocation and reclamation of paging space slots** in the *Performance management and tuning*.

The **Understanding Memory Mapping** in *AIX 5L Version 5.3 General Programming Concepts*.

---

## dlclose Subroutine

### Purpose

Closes and unloads a module loaded by the **dlopen** subroutine.

### Syntax

```
#include <dlfcn.h>
int dlclosel(Data);
void *Data;
```

## Description

The **dlclose** subroutine is used to remove access to a module loaded with the **dlopen** subroutine. In addition, access to dependent modules of the module being unloaded is removed as well.

The **dlclose** subroutine performs C++ termination, like the **terminateAndUnload** subroutine does.

Modules being unloaded with the **dlclose** subroutine will not be removed from the process's address space if they are still required by other modules. Nevertheless, subsequent uses of *Data* are invalid, and further uses of symbols that were exported by the module being unloaded result in undefined behavior.

## Parameters

*Data*            A loaded module reference returned from a previous call to **dlopen**.

## Return Values

Upon successful completion, 0 (zero) is returned. Otherwise, **errno** is set to **EINVAL**, and the return value is also **EINVAL**. Even if the **dlclose** subroutine succeeds, the specified module may still be part of the process's address space if the module is still needed by other modules.

## Error Codes

**EINVAL**            The *Data* parameter does not refer to a module opened by **dlopen** that is still open. The parameter may be corrupt or the module may have been unloaded by a previous call to **dlclose**.

## Related Information

The **dlerror** (“dlerror Subroutine”) subroutine, **dlopen** (“dlopen Subroutine” on page 217) subroutine, **load** (“load and loadAndInit Subroutines” on page 727) subroutine, **loadquery** (“loadquery Subroutine” on page 733) subroutine, **unload** subroutine, **loadbind** (“loadbind Subroutine” on page 731) subroutine.

The **ld** command.

The Shared Libraries and Shared Memory Overview and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## dlerror Subroutine

### Purpose

Returns a pointer to information about the last **dlopen**, **dlsym**, or **dlclose** error.

### Syntax

```
#include <dlfcn.h>
char *dlerror(void);
```

### Description

The **dlerror** subroutine is used to obtain information about the last error that occurred in a dynamic loading routine (that is, **dlopen**, **dlsym**, or **dlclose**). The returned value is a pointer to a null-terminated string without a final newline. Once a call is made to this function, subsequent calls without any intervening dynamic loading errors will return NULL.



Applications can avoid calling the **dlderror** subroutine, in many cases, by examining **errno** after a failed call to a dynamic loading routine. If **errno** is **ENOEXEC**, the **dlderror** subroutine will return additional information. In all other cases, **dlderror** will return the string corresponding to the value of **errno**.

The **dlderror** function may invoke **loadquery** to ascertain reasons for a failure. If a call is made to **load** or **unload** between calls to **dlopen** and **dlderror**, incorrect information may be returned.

## Return Values

A pointer to a static buffer is returned; a NULL value is returned if there has been no error since the last call to **dlderror**. Applications should not write to this buffer; they should make a copy of the buffer if they wish to preserve the buffer's contents.

## Related Information

The **load** ("load and loadAndInit Subroutines" on page 727) subroutine, **loadbind** ("loadbind Subroutine" on page 731) subroutine, **loadquery** ("loadquery Subroutine" on page 733) subroutine, **unload** subroutine, **dlopen** ("dlopen Subroutine") subroutine, **dlclose** ("dlclose Subroutine" on page 215) subroutine.

The **ld** command.

The Shared Libraries and Shared Memory Overview and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## dlopen Subroutine

### Purpose

Dynamically loads a module into the calling process.

### Syntax

```
#include <dlfcn.h>
void *dlopen (FilePath, Flags);
const char *FilePath;
int Flags;
```

### Description

The **dlopen** subroutine loads the module specified by *FilePath* into the executing process's address space. Dependents of the module are automatically loaded as well. If the module is already loaded, it is not loaded again, but a new, unique value will be returned by the **dlopen** subroutine.

The **dlopen** subroutine is a portable way of dynamically loading shared libraries. It performs C++ static initialization of the modules that it loads, like the **loadAndInit** subroutine does.

The value returned by the **dlopen** might be used in subsequent calls to **dlsym** and **dlclose**. If an error occurs during the operation, **dlopen** returns NULL.

If the main application was linked with the **-brtl** option, then the runtime linker is invoked by **dlopen**. If the module being loaded was linked with runtime linking enabled, both intra-module and inter-module references are overridden by any symbols available in the main application. If runtime linking was enabled, but the module was not built enabled, then all inter-module references will be overridden, but some intra-module references will not be overridden.

If the module being opened with **dlopen** or any of its dependents is being loaded for the first time, initialization routines for these newly-loaded routines are called (after runtime linking, if applicable) before

**dlopen** returns. Initialization routines are the functions specified with the **-binitfni**: linker option when the module was built. (See the **ld** command for more information about this option.)

The modules that have the **binitfni** functions are initialized first, regardless of the priority that is specified using the **-qmkshrobj** option. That is to say, if the modules have both the **binitfni** functions and the **qmkshrobj** priorities specified, the **binitfni** functions are initialized first.

If a **dlopen** subroutine is called from within a **binitfni** function, the initialization of the current module is abandoned for other modules.

#### Notes:

1. The initialization functions need not have any special names, and multiple functions per module are allowed.
2. If the module being loaded has read-other permission, the module is loaded into the global shared library segment. Modules loaded into the global shared library segment are not unloaded even if they are no longer being used. Use the **slibclean** command to remove unused modules from the global shared library segment.

The **LIBPATH** or **LD\_LIBRARY\_PATH** environment variables can be used to specify a list of directories in which **dlopen** searches for the named module. The running application also contains a set of library search paths that were specified when the application was linked; these paths are searched after any paths found in **LIBPATH** or **LD\_LIBRARY\_PATH**.

*FilePath* Specifies the name of a file containing the loadable module. This parameter can be contain an absolute path, a relative path, or no path component. If *FilePath* contains a slash character, *FilePath* is used directly, and no directories are searched.

If the *FilePath* parameter is `/unix`, **dlopen** returns a value that can be used to look up symbols in the current kernel image, including those symbols found in any kernel extension that was available at the time the process began execution.

If the value of *FilePath* is `NULL`, a value for the main application is returned. This allows dynamically loaded objects to look up symbols in the main executable, or for an application to examine symbols available within itself.

## Flags

Specifies variations of the behavior of **dlopen**. Either **RTLD\_NOW** or **RTLD\_LAZY** must always be specified. Other flags may be OR'ed with **RTLD\_NOW** or **RTLD\_LAZY**.

<b>RTLD_NOW</b>	Load all dependents of the module being loaded and resolve all symbols.
<b>RTLD_LAZY</b>	Specifies the same behavior as <b>RTLD_NOW</b> . In a future release of the operating system, the behavior of the <b>RTLD_LAZY</b> may change so that loading of dependent modules is deferred of resolution of some symbols is deferred.
<b>RTLD_GLOBAL</b>	Allows symbols in the module being loaded to be visible when resolving symbols used by other <b>dlopen</b> calls. These symbols will also be visible when the main application is opened with <b>dlopen</b> ( <code>NULL</code> , <i>mode</i> ).
<b>RTLD_LOCAL</b>	Prevent symbols in the module being loaded from being used when resolving symbols used by other <b>dlopen</b> calls. Symbols in the module being loaded can only be accessed by calling <b>dlsym</b> subroutine. If neither <b>RTLD_GLOBAL</b> nor <b>RTLD_LOCAL</b> is specified, the default is <b>RTLD_LOCAL</b> . If both flags are specified, <b>RTLD_LOCAL</b> is ignored.
<b>RTLD_MEMBER</b>	The <b>dlopen</b> subroutine can be used to load a module that is a member of an archive. The <b>L_LOADMEMBER</b> flag is used when the <b>load</b> subroutine is called. The module name <i>FilePath</i> names the archive and archive member according to the rules outlined in the <b>load</b> subroutine.

## RTLD\_NOAUTODEFER

Prevents deferred imports in the module being loaded from being automatically resolved by subsequent loads. The **L\_NOAUTODEFER** flag is used when the **load** subroutine is called.

Ordinarily, modules built for use by the **dlopen** and **dlsym** sub routines will not contain deferred imports. However, deferred imports can be still used. A module opened with **dlopen** may provide definitions for deferred imports in the main application, for modules loaded with the **load** subroutine (if the **L\_NOAUTODEFER** flag was not used), and for other modules loaded with the **dlopen** subroutine (if the **RTLD\_NOAUTODEFER** flag was not used).

## Return Values

Upon successful completion, **dlopen** returns a value that can be used in calls to the **dlsym** and **dlclose** subroutines. The value is not valid for use with the **loadbind** and **unload** subroutines.

If the **dlopen** call fails, NULL (a value of 0) is returned and the global variable **errno** is set. If **errno** contains the value ENOEXEC, further information is available via the **dlerror** function.

## Error Codes

See the **load** subroutine for a list of possible **errno** values and their meanings.

## Related Information

The **dlclose** (“dlclose Subroutine” on page 215) subroutine, **dlerror** (“dlerror Subroutine” on page 216) subroutine, **load** (“load and loadAndInit Subroutines” on page 727) subroutine, **loadbind** (“loadbind Subroutine” on page 731) subroutine, **loadquery** (“loadquery Subroutine” on page 733) subroutine, **unload** subroutine.

The **ld** command.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

The Dynamically loading a shared library section in the *XL C/C++ V8.0 for AIX Programming Guide* book.

---

## dlsym Subroutine

### Purpose

Looks up the location of a symbol in a module that is loaded with **dlopen**.

### Syntax

```
#include <dlfcn.h>
void *dlsym(Handle, Symbol);
void *Handle;
const char *Symbol;
```

### Description

The **dlsym** subroutine looks up a named symbol exported from a module loaded by a previous call to the **dlopen** subroutine. Only exported symbols are found by **dlsym**. See the **ld** command to see how to export symbols from a module.

*Handle* Specifies a value returned by a previous call to **dlopen** or one of the special handles **RTLD\_DEFAULT**, **RTLD\_NEXT** or **RTLD\_MYSELF**.

*Symbol* Specifies the name of a symbol exported from the referenced module in the form of a NULL-terminated string or the special symbol name **RTLD\_ENTRY**.

**Note:** C++ symbol names should be passed to **dlsym** in mangled form; **dlsym** does not perform any name demangling on behalf of the calling application.

In case of the special handle **RTLD\_DEFAULT**, **dlsym** searches for the named symbol starting with the first module loaded. It then proceeds through the list of initial loaded modules and any global modules obtained with **dlopen** until a match is found. This search follows the default model employed to relocate all modules within the process.

In case of the special handle **RTLD\_NEXT**, **dlsym** searches for the named symbol in the modules that were loaded following the module from which the **dlsym** call is being made.

In case of the special handle **RTLD\_MYSELF**, **dlsym** searches for the named symbol in the modules that were loaded starting with the module from which the **dlsym** call is being made.

In case of the special symbol name **RTLD\_ENTRY**, **dlsym** returns the module's entry point. The entry point, if present, is the value of the module's loader section symbol marked as entry point.

In case of **RTLD\_DEFAULT**, **RTLD\_NEXT**, and **RTLD\_MYSELF**, if the modules being searched have been loaded from **dlopen** calls, **dlsym** searches the module only if the caller is part of the same **dlopen** dependency hierarchy, or if the module was given global search access. See **dlopen** for a discussion of the **RTLD\_GLOBAL** mode.

A search for the named symbol is based upon breadth-first ordering of the module and its dependants. If the module was constructed using the **-G** or **-brtl** linker option, the module's dependants will include all modules named on the **ld** command line, in the original order. The dependants of a module that was not linked with the **-G** or **-brtl** linker option will be listed in an unspecified order.

## Return Values

If the named symbol is found, its address is returned. If the named symbol is not found, NULL is returned and **errno** is set to 0. If *Handle* or *Symbol* is invalid, NULL is returned and **errno** is set to **EINVAL**.

If the first definition found is an export of an imported symbol, this definition will satisfy the search. The address of the imported symbol is returned. If the first definition is a deferred import, the definition is ignored and the search continues.

If the named symbol refers to a BSS symbol (uninitialized data structure), the search continues until an initialized instance of the symbol is found or the module and all of its dependants have been searched. If an initialized instance is found, its address is returned; otherwise, the address of the first uninitialized instance is returned.

## Error Codes

**EINVAL** If the *Handle* parameter does not refer to a module opened by **dlopen** that is still loaded or if the *Symbol* parameter points to an invalid address, the **dlsym** subroutine returns NULL and **errno** is set to **EINVAL**.

## Related Information

The **dlclose** (“dlclose Subroutine” on page 215) subroutine, **dlerror** (“dlerror Subroutine” on page 216) subroutine, **dlopen** (“dlopen Subroutine” on page 217) subroutine, **load** (“load and loadAndInit

Subroutines” on page 727) subroutine, **loadbind** (“loadbind Subroutine” on page 731) subroutine, **loadquery** (“loadquery Subroutine” on page 733) subroutine, **unload** subroutine.

The **ld** command.

---

## **drand48, erand48, jrand48, lcong48, lrand48, mrand48, nrand48, seed48, or srand48 Subroutine**

### **Purpose**

Generate uniformly distributed pseudo-random number sequences.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

```
#include <stdlib.h>
double drand48 (void)

double erand48 ( xsubi)
unsigned short int xsubi[3];

long int jrand48 (xsubi)
unsigned short int xsubi[3];

void lcong48 ( Parameter)
unsigned short int Parameter[7];

long int lrand48 (void)
long int mrand48 (void)
long int nrand48 (xsubi)
unsigned short int xsubi[3];

unsigned short int *seed48 ( Seed16v)
unsigned short int Seed16v[3];

void srand48 ( SeedValue)
long int SeedValue;
```

### **Description**

**Attention:** Do not use the **drand48, erand48, jrand48, lcong48, lrand48, mrand48, nrand48, seed48, or srand48** subroutine in a multithreaded environment.

This family of subroutines generates pseudo-random numbers using the linear congruential algorithm and 48-bit integer arithmetic.

The **drand48** subroutine and the **erand48** subroutine return positive double-precision floating-point values uniformly distributed over the interval [0.0, 1.0).

The **lrand48** subroutine and the **nrand48** subroutine return positive long integers uniformly distributed over the interval [0,2\*\*31).

The **mrnd48** subroutine and the **jrnd48** subroutine return signed long integers uniformly distributed over the interval [-2\*\*31, 2\*\*31).

The **srand48** subroutine, **seed48** subroutine, and **lcong48** subroutine initialize the random-number generator. Programs must call one of them before calling the **drand48**, **lrand48** or **mrand48** subroutines. (Although it is not recommended, constant default initializer values are supplied if the **drand48**, **lrand48** or **mrand48** subroutines are called without first calling an initialization subroutine.) The **erand48**, **nrand48**, and **jrand48** subroutines do not require that an initialization subroutine be called first.

The previous value pointed to by the **seed48** subroutine is stored in a 48-bit internal buffer, and a pointer to the buffer is returned by the **seed48** subroutine. This pointer can be ignored if it is not needed, or it can be used to allow a program to restart from a given point at a later time. In this case, the pointer is accessed to retrieve and store the last value pointed to by the **seed48** subroutine, and this value is then used to reinitialize, by means of the **seed48** subroutine, when the program is restarted.

All the subroutines work by generating a sequence of 48-bit integer values,  $x[i]$ , according to the linear congruential formula:

$$x[n+1] = (ax[n] + c) \bmod m, \quad n \text{ is } \geq 0$$

The parameter  $m = 248$ ; hence 48-bit integer arithmetic is performed. Unless the **lcong48** subroutine has been called, the multiplier value  $a$  and the addend value  $c$  are:

$a = 5DEECE66D$  base 16 = 273673163155 base 8

$c = B$  base 16 = 13 base 8

## Parameters

<i>xsubi</i>	Specifies an array of three shorts, which, when concatenated together, form a 48-bit integer.
<i>SeedValue</i>	Specifies the initialization value to begin randomization. Changing this value changes the randomization pattern.
<i>Seed16v</i>	Specifies another seed value; an array of three unsigned shorts that form a 48-bit seed value.
<i>Parameter</i>	Specifies an array of seven shorts, which specifies the initial <i>xsubi</i> value, the multiplier value $a$ and the add-in value $c$ .

## Return Values

The value returned by the **drand48**, **erand48**, **jrand48**, **lrand48**, **nrand48**, and **mrand48** subroutines is computed by first generating the next 48-bit  $x[i]$  in the sequence. Then the appropriate number of bits, according to the type of data item to be returned, are copied from the high-order (most significant) bits of  $x[i]$  and transformed into the returned value.

The **drand48**, **lrand48**, and **mrand48** subroutines store the last 48-bit  $x[i]$  generated into an internal buffer; this is why they must be initialized prior to being invoked.

The **erand48**, **jrand48**, and **nrand48** subroutines require the calling program to provide storage for the successive  $x[i]$  values in the array pointed to by the *xsubi* parameter. This is why these routines do not have to be initialized; the calling program places the desired initial value of  $x[i]$  into the array and pass it as a parameter.

By using different parameters, the **erand48**, **jrand48**, and **nrand48** subroutines allow separate modules of a large program to generate independent sequences of pseudo-random numbers. In other words, the sequence of numbers that one module generates does not depend upon how many times the subroutines are called by other modules.

The **lcong48** subroutine specifies the initial  $x[i]$  value, the multiplier value  $a$ , and the addend value  $c$ . The *Parameter* array elements *Parameter*[0-2] specify  $x[i]$ , *Parameter*[3-5] specify the multiplier  $a$ , and *Parameter*[6] specifies the 16-bit addend  $c$ . After **lcong48** has been called, a subsequent call to either the **srand48** or **seed48** subroutine restores the standard  $a$  and  $c$  specified before.

The initializer subroutine **seed48** sets the value of  $x[i]$  to the 48-bit value specified in the array pointed to by the *Seed16v* parameter. In addition, **seed48** returns a pointer to a 48-bit internal buffer that contains the previous value of  $x[i]$  that is used only by **seed48**. The returned pointer allows you to restart the pseudo-random sequence at a given point. Use the pointer to copy the previous  $x[i]$  value into a temporary array. Then call **seed48** with a pointer to this array to resume processing where the original sequence stopped.

The initializer subroutine **srand48** sets the high-order 32 bits of  $x[i]$  to the 32 bits contained in its parameter. The low order 16 bits of  $x[i]$  are set to the arbitrary value 330E16.

## Related Information

The **rand**, **srand** subroutine, **random**, **srandom**, **initstate**, or **setstate** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## drem Subroutine

### Purpose

Computes the IEEE Remainder as defined in the IEEE Floating-Point Standard.

### Libraries

IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
```

```
double drem ( x, y)  
double x, y;
```

### Description

The **drem** subroutine calculates the remainder  $r$  equal to  $x$  minus  $n$  to the  $x$  power multiplied by  $y$  ( $r = x - n * y$ ), where the  $n$  parameter is the integer nearest the exact value of  $x$  divided by  $y$  ( $x/y$ ). If  $|n - x/y| = 1/2$ , then the  $n$  parameter is an even value. Therefore, the remainder is computed exactly, and the absolute value of  $r$  ( $|r|$ ) is less than or equal to the absolute value of  $y$  divided by 2 ( $|y|/2$ ).

The IEEE Remainder differs from the **fmod** subroutine in that the IEEE Remainder always returns an  $r$  parameter such that  $|r|$  is less than or equal to  $|y|/2$ , while **FMOD** returns an  $r$  such that  $|r|$  is less than or equal to  $|y|$ . The IEEE Remainder is useful for argument reduction for transcendental functions.

**Note:** Compile any routine that uses subroutines from the **libm.a** library with the **-lm** flag. For example: compile the **drem.c** file:

```
cc drem.c -lm
```

**Note:** For new development, the **remainder** subroutine is the preferred interface.

### Parameters

$x$  Specifies double-precision floating-point value.  
 $y$  Specifies a double-precision floating-point value.

## Return Values

The **drem** subroutine returns a NaNQ value for (x, 0) and (+/-INF, y).

## Related Information

The **floor**, **ceil**, **nearest**, **trunc**, **rint**, **itrunc**, **fmod**, **fabs**, or **uitrunc** (“floor, floorf, floorl, floord32, floord64, floord128, nearest, trunc, itrunc, and uitrunc Subroutines” on page 277) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## **\_end**, **\_etext**, or **\_edata** Identifier

### Purpose

Define the first addresses following the program, initialized data, and all data.

### Syntax

```
extern _end;  
extern _etext;  
extern _edata;
```

### Description

The external names **\_end**, **\_etext**, and **\_edata** are defined by the loader for all programs. They are not subroutines but identifiers associated with the following addresses:

<b>_etext</b>	The first address following the program text.
<b>_edata</b>	The first address following the initialized data region.
<b>_end</b>	The first address following the data region that is not initialized. The name <b>end</b> (with no underscore) defines the same address as does <b>_end</b> (with underscore).

The break value of the program is the first location beyond the data. When a program begins running, this location coincides with **end**. However, many factors can change the break value, including:

- The **brk** or **sbrk** subroutine
- The **malloc** subroutine
- The standard I/O subroutines
- The **-p** flag with the **cc** command

Therefore, use the **brk** or **sbrk(0)** subroutine, not the **end** address, to determine the break value of the program.

### Related Information

The **brk** or **sbrk** (“brk or sbrk Subroutine” on page 122) subroutine, **malloc** (“malloc, free, realloc, calloc, mallopt, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.



---

## ecvt, fcvt, or gcvt Subroutine

### Purpose

Converts a floating-point number to a string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
char *ecvt ( Value, NumberOfDigits, DecimalPointer, Sign;)
```

```
double Value;
```

```
int NumberOfDigits, *DecimalPointer, *Sign;
```

```
char *fcvt (Value, NumberOfDigits, DecimalPointer, Sign;)
```

```
double Value;
```

```
int NumberOfDigits, *DecimalPointer, *Sign;
```

```
char *gcvt (Value, NumberOfDigits, Buffer;)
```

```
double Value;
```

```
int NumberOfDigits;
```

```
char *Buffer;
```

### Description

The **ecvt**, **fcvt**, and **gcvt** subroutines convert floating-point numbers to strings.

The **ecvt** subroutine converts the *Value* parameter to a null-terminated string and returns a pointer to it. The *NumberOfDigits* parameter specifies the number of digits in the string. The low-order digit is rounded according to the current rounding mode. The **ecvt** subroutine sets the integer pointed to by the *DecimalPointer* parameter to the position of the decimal point relative to the beginning of the string. (A negative number means the decimal point is to the left of the digits given in the string.) The decimal point itself is not included in the string. The **ecvt** subroutine also sets the integer pointed to by the *Sign* parameter to a nonzero value if the *Value* parameter is negative and sets a value of 0 otherwise.

The **fcvt** subroutine operates identically to the **ecvt** subroutine, except that the correct digit is rounded for C or FORTRAN F-format output of the number of digits specified by the *NumberOfDigits* parameter.

**Note:** In the F-format, the *NumberOfDigits* parameter is the number of digits desired after the decimal point. Large numbers produce a long string of digits before the decimal point, and then *NumberOfDigits* digits after the decimal point. Generally, the **gcvt** and **ecvt** subroutines are more useful for large numbers.

The **gcvt** subroutine converts the *Value* parameter to a null-terminated string, stores it in the array pointed to by the *Buffer* parameter, and then returns the *Buffer* parameter. The **gcvt** subroutine attempts to produce a string of the *NumberOfDigits* parameter significant digits in FORTRAN F-format. If this is not possible, the E-format is used. The **gcvt** subroutine suppresses trailing zeros. The string is ready for printing, complete with minus sign, decimal point, or exponent, as appropriate. The radix character is determined by the current locale (see **setlocale** subroutine). If the **setlocale** subroutine has not been called successfully, the default locale, POSIX, is used. The default locale specifies a . (period) as the radix character. The **LC\_NUMERIC** category determines the value of the radix character within the current locale.

The **ecvt**, **fcvt**, and **gcvt** subroutines represent the following special values that are specified in ANSI/IEEE standards 754-1985 and 854-1987 for floating-point arithmetic:

<b>Quiet NaN</b>	Indicates a quiet not-a-number (NaNQ)
<b>Signalling NaN</b>	Indicates a signaling NaNs
<b>Infinity</b>	Indicates a INF value

The sign associated with each of these values is stored in the *Sign* parameter.

**Note:** A value of 0 can be positive or negative. In the IEEE floating-point, zeros also have signs and set the *Sign* parameter appropriately.

**Attention:** All three subroutines store the strings in a static area of memory whose contents are overwritten each time one of the subroutines is called.

## Parameters

<i>Value</i>	Specifies some double-precision floating-point value.
<i>NumberOfDigits</i>	Specifies the number of digits in the string.
<i>DecimalPointer</i>	Specifies the position of the decimal point relative to the beginning of the string.
<i>Sign</i>	Specifies that the sign associated with the return value is placed in the <i>Sign</i> parameter. In IEEE floating-point, since 0 can be signed, the <i>Sign</i> parameter is set appropriately for signed 0.
<i>Buffer</i>	Specifies a character array for the string.

## Related Information

The **atof**, **strtod**, **atoff**, or **strtodf** (“atof atoff Subroutine” on page 96) subroutine, **fp\_read\_rnd**, or **fp\_swap\_rnd** (“fp\_read\_rnd or fp\_swap\_rnd Subroutine” on page 302) subroutine, **printf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **scanf** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## EnableCriticalSection, BeginCriticalSection, and EndCriticalSection Subroutine

### Purpose

Enables a thread to be exempted from timeslicing and signal suspension, and protects critical sections.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/thread_ct1.h>
```

```
int EnableCriticalSection(void);  
void BeginCriticalSection(void);  
void EndCriticalSection(void);
```

## Description

When called, the **EnableCriticalSections** subroutine enables the thread to be exempted from timeslicing and signal suspension. Once that is done, the thread can call the **BeginCriticalSection** and **EndCriticalSection** subroutines to protect critical sections. Calling the **BeginCriticalSection** and **EndCriticalSection** subroutines with exemption disabled has no effect. The subroutines are safe for use by multithreaded applications.

Once the service is enabled, the thread can protect critical sections by calling the **BeginCriticalSection** and **EndCriticalSection** subroutines. Calling the **BeginCriticalSection** subroutine will exempt the thread from timeslicing and suspension. Calling the **EndCriticalSection** subroutine will clear exemption for the thread.

The **BeginCriticalSection** subroutine will not make a system call. The **EndCriticalSection** subroutine might make a system call if the thread was granted a benefit during the critical section. The purpose of the system call would be to notify the kernel that any posted but undelivered stop signals can be delivered, and any postponed timeslice can now be completed.

## Return Values

The **EnableCriticalSections** subroutine returns a zero.

---

## erf, erff, or erfl Subroutine

### Purpose

Computes the error and complementary error functions.

### Libraries

IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
```

```
double erf ( x )
double x;
float erff ( x )
float x;
long double erfl ( x )
long double x;
```

### Description

The **erf**, **erff**, and **erfl** subroutines return the error function of the  $x$  parameter, defined for the **erf** subroutine as the following:

$$\text{erf}(x) = (2/\sqrt{\pi}) * (\text{integral } [0 \text{ to } x] \text{ of } \exp(-t^2)) dt$$

$$\text{erfc}(x) = 1.0 - \text{erf}(x)$$

**Note:** Compile any routine that uses subroutines from the **libm.a** library with the **-lm** flag. To compile the **erf.c** file, for example, enter:

```
cc erf.c -lm
```

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies a double-precision floating-point value.

## Return Values

Upon successful completion, the **erf**, **erff**, and **erfl** subroutines return the value of the error function.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ ,  $\pm 0$  is returned.

If *x* is  $\pm \text{Inf}$ ,  $\pm 1$  is returned.

If *x* is subnormal, a range error may occur, and  $2 * x/\text{sqrt}(\pi)$  should be returned.

## Related Information

“erfc, erfcf, or erfcl Subroutine,” “exp, expf, or expl Subroutine” on page 245, “feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

The sqrt, sqrtf, or sqrtl Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## erfc, erfcf, or erfcl Subroutine

### Purpose

Computes the complementary error function.

### Syntax

```
#include <math.h>
```

```
float erfcf (x)  
float x;
```

```
long double erfcl (x)  
long double x;
```

```
double erfc (x)  
double x;
```

## Description

The **erfcf**, **erfcl**, and **erfc** subroutines compute the complementary error function  $1.0 - \text{erf}(x)$ .

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **erfcf**, **erfcl**, and **erfc** subroutines return the value of the complementary error function.

If the correct value would cause underflow and is not representable, a range error may occur. Either 0.0 (if representable), or an implementation-defined value is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ , +1 is returned.

If *x* is  $-\text{Inf}$ , +2 is returned.

If *x* is  $+\text{Inf}$ , +0 is returned.

If the correct value would cause underflow and is representable, a range error may occur and the correct value is returned.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## errlog Subroutine

### Purpose

Logs an application error to the system error log.

### Library

Run-Time Services Library (**librts.a**)

### Syntax

```
#include <sys/errids.h>
int errlog ( ErrorStructure, Length)
void *ErrorStructure;
unsigned int Length;
```

## Description

The **errlog** subroutine writes an error log entry to the **/dev/error** file. The **errlog** subroutine is used by application programs.

The transfer from the **err\_rec** structure to the error log is by a **write** subroutine to the **/dev/error** special file.

The **errdemon** process reads from the **/dev/error** file and writes the error log entry to the system error log. The timestamp, machine ID, node ID, and Software Vital Product Data associated with the resource name (if any) are added to the entry before going to the log.

## Parameters

### *ErrorStructure*

Points to an error record structure containing an error record. Valid error record structures are typed in the **/usr/include/sys/err\_rec.h** file. The two error record structures available are **err\_rec** and **err\_rec0**. The **err\_rec** structure is used when the **detail\_data** field is required. When the **detail\_data** field is not required, the **err\_rec0** structure is used.

```
struct err_rec0 {
    unsigned int error_id;
    char resource_name[ERR_NAMESIZE];
};
struct err_rec {
    unsigned int error_id;
    char resource_name[ERR_NAMESIZE];
    char detail_data[1];
};
```

The fields of the structures **err\_rec** and **err\_rec0** are:

#### **error\_id**

Specifies an index for the system error template database, and is assigned by the **errupdate** command when adding an error template. Use the **errupdate** command with the **-h** flag to get a **#define** statement for this 8-digit hexadecimal index.

#### **resource\_name**

Specifies the name of the resource that has detected the error. For software errors, this is the name of a software component or an executable program. For hardware errors, this is the name of a device or system component. It does not indicate that the component is faulty or needs replacement instead, it is used to determine the appropriate diagnostic modules to be used to analyze the error.

#### **detail\_data**

Specifies an array from 0 to **ERR\_REC\_MAX** bytes of user-supplied data. This data may be displayed by the **errpt** command in hexadecimal, alphanumeric, or binary form, according to the **data\_encoding** fields in the error log template for this **error\_id** field.

### *Length*

Specifies the length in bytes of the **err\_rec** structure, which is equal to the size of the **error\_id** and **resource\_name** fields plus the length in bytes of the **detail\_data** field.

## Return Values

- 0 The entry was logged successfully.
- 1 The entry was not logged.

## Files

<code>/dev/error</code>	Provides standard device driver interfaces required by the error log component.
<code>/usr/include/sys/errids.h</code>	Contains definitions for error IDs.
<code>/usr/include/sys/err_rec.h</code>	Contains structures defined as arguments to the <b>errsave</b> kernel service and the <b>errlog</b> subroutine.
<code>/var/adm/ras/errlog</code>	Maintains the system error log.

## Related Information

The **errclear**, **errdead**, **errinstall**, **errlogger**, **errmsg**, **errpt**, **errstop**, and **errupdate** commands.

The **errlog\_open**, **errlog\_close**, **errlog\_find\_first**, **errlog\_find\_next**, **errlog\_find\_sequence**, **errlog\_set\_direction**, and **errlog\_write** subroutines.

The `/dev/error` special file.

The **errdemon** daemon.

The **errsave** kernel service.

Error Logging Overview in *Messages Guide and Reference*.

---

## errlog\_close Subroutine

### Purpose

Closes an open error log file.

### Syntax

```
library liberrlog.a

#include <sys/errlog.h>

int errlog_close(handle)
errlog_handle_t handle;
```

### Description

The error log specified by the `handle` argument is closed. The `handle` must have been returned from a previous **errlog\_open** call.

### Return Values

Upon successful completion, the **errlog\_close** subroutine returns 0.

If an error occurs, the **errlog\_close** subroutine returns **LE\_ERR\_INVARG**.

## Related Information

The **errlog\_open**, **errlog\_find\_first**, **errlog\_find\_next**, **errlog\_find\_sequence**, **errlog\_set\_direction**, **errlog\_write**, and **errlog** subroutines.

---

## errlog\_find\_first, errlog\_find\_next, and errlog\_find\_sequence Subroutines

### Purpose

Retrieves an error log entry using supplied criteria.

### Syntax

library liberrlog.a

```
#include <sys/errlog.h>
```

```
int errlog_find_first(handle, filter, result)
errlog_handle_t handle;
errlog_match_t *filter;
errlog_entry_t *result;
```

```
int errlog_find_next(handle, result)
errlog_handle_t handle;
errlog_entry_t *result;
```

```
int errlog_find_sequence(handle, sequence, result)
errlog_handle_t handle;
int sequence;
errlog_entry_t *result;
```

### Description

The **errlog\_find\_first** subroutine finds the first occurrence of the search argument specified by filter using the direction specified by the **errlog\_set\_direction** subroutine. The reverse direction is used if none was specified. In other words, by default, entries are searched starting with the most recent entry.

The **errlog\_match\_t** structure, pointed to by the filter parameter, defines a test expression or set of expressions to be applied to each errlog entry.

If the value passed in the filter parameter is null, the **errlog\_find\_first** subroutine returns the first entry in the log, and the **errlog\_find\_next** subroutine can then be used to return subsequent entries. To read all log entries in the desired direction, open the log, then issue **errlog\_find\_next** calls.

To define a basic expression, **em\_field** must be set to the field in the errlog entry to be tested, **em\_op** must be set to the relational operator to be applied to that field, and either **em\_intvalue** or **em\_strvalue** must be set to the value to test against. Basic expressions may be combined by attaching them to **em\_left** and **em\_right** of another **errlog\_match\_t** structure and setting **em\_op** of that structure to a binary or unary operator. These complex expressions may then be combined with other basic or complex expressions in the same fashion to build a tree that can define a filter of arbitrary complexity.

The **errlog\_find\_next** subroutine finds the next error log entry matching the criteria specified by a previous **errlog\_find\_first** call. The search continues in the direction specified by the **errlog\_set\_direction** subroutine or the reverse direction by default.

The **errlog\_find\_sequence** subroutine returns the entry matching the specified error log sequence number, found in the **el\_sequence** field of the **errlog\_entry** structure.

### Parameters

The handle contains the handle returned by a prior call to **errlog\_open**.

The filter parameter points to an **errlog\_match\_t** element defining the search argument, or the first of an argument tree.



The sequence parameter contains the sequence number of the entry to be retrieved.

The result parameter must point to the area to contain the returned error log entry.

## Return Values

Upon successful completion, the **errlog\_find\_first**, **errlog\_find\_next**, and **errlog\_find\_sequence** subroutines return 0, and the memory referenced by result contains the found entry.

The following errors may be returned:

<b>LE_ERR_INVARG</b>	A parameter error was detected.
<b>LE_ERR_NOMEM</b>	Memory could not be allocated.
<b>LE_ERR_IO</b>	An i/o error occurred.
<b>LE_ERR_DONE</b>	No more entries were found.

## Examples

The code below demonstrates how to search for all errlog entries in a date range and with a class of **H** (hardware) or **S** (software).

```
{
    extern int          begintime, endtime;

    errlog_match_t     beginstamp, endstamp, andstamp;
    errlog_match_t     hardclass, softclass, orclass;
    errlog_match_t     andtop;
    int                ret;
    errlog_entry_t     result;

    /*
     * Select begin and end times
     */
    beginstamp.em_op = LE_OP_GT;           /* Expression 'A' */
    beginstamp.em_field = LE_MATCH_TIMESTAMP;
    beginstamp.em_intvalue = begintime;

    endstamp.em_op = LE_OP_LT;           /* Expression 'B' */
    endstamp.em_field = LE_MATCH_TIMESTAMP;
    endstamp.em_intvalue = endtime;

    andstamp.em_op = LE_OP_AND;          /* 'A' and 'B' */
    andstamp.em_left = &beginstamp;
    andstamp.em_right = &endstamp;

    /*
     * Select the classes we're interested in.
     */
    hardclass.em_op = LE_OP_EQ;          /* Expression 'C' */
    hardclass.em_field = LE_MATCH_CLASS;
    hardclass.em_strvalue = "H";

    softclass.em_op = LE_OP_EQ;         /* Expression 'D' */
    softclass.em_field = LE_MATCH_CLASS;
    softclass.em_strvalue = "S";

    orclass.em_op = LE_OP_OR;           /* 'C' or 'D' */
    orclass.em_left = &hardclass;
    orclass.em_right = &softclass;

    andtop.em_op = LE_OP_AND;           /* ('A' and 'B') and ('C' or 'D') */
    andtop.em_left = &andstamp;
```

```

    andtop.em_right = &orclass;

    ret = errlog_find_first(handle, &andtop, &result);
}

```

The **errlog\_find\_first** function will return the first entry matching filter. Successive calls to the **errlog\_find\_next** function will return successive entries that match the filter specified in the most recent call to the **errlog\_find\_first** function. When no more matching entries are found, the **errlog\_find\_first** and **errlog\_find\_next** functions will return the value **LE\_ERR\_DONE**.

## Related Information

The **errlog\_open**, **errlog\_close**, **errlog\_set\_direction**, **errlog\_write**, and **errlog** subroutines.

---

## errlog\_open Subroutine

### Purpose

Opens an error log and returns a handle for use with other **liberrlog.a** functions.

### Syntax

library liberrlog.a

```

#include <fcntl.h>
#include <sys/errlog.h>

```

```

int errlog_open(path, mode, magic, handle)
char          *path;
int           mode;
unsigned int  magic;
errlog_handle_t *handle;

```

### Description

The error log specified by the path argument will be opened using mode. The handle pointed to by the handle parameter must be used with subsequent operations.

### Parameters

The path parameter specifies the path to the log file to be opened. If path is NULL, the default errlog file will be opened. The valid values for mode are the same as they are for the open system subroutine. They can be found in the **fcntl.h** files.

The **magic** argument takes the **LE\_MAGIC** value, indicating which version of the **errlog\_entry\_t** structure this application was compiled with.

### Return Values

Upon successful completion, the **errlog\_open** subroutine returns a 0 and sets the memory pointed to by handle to a handle used by subsequent **liberrlog** operations.

Upon error, the **errlog\_open** subroutine returns one of the following:

<b>LE_ERR_INVARG</b>	A parameter error was detected.
<b>LE_ERR_NOFILE</b>	The log file does not exist.
<b>LE_ERR_NOMEM</b>	Memory could not be allocated.
<b>LE_ERR_IO</b>	An i/o error occurred.
<b>LE_ERR_INVFILE</b>	The file is not a valid error log.

## Related Information

The `errlog_close`, `errlog_find_first`, `errlog_find_next`, `errlog_find_sequence`, `errlog_set_direction`, `errlog_write`, and `errlog` subroutines.

The `/usr/include/fcntl.h` include files found in *AIX 5L Version 5.3 Files Reference*.

---

## errlog\_set\_direction Subroutine

### Purpose

Sets the direction for the error log find functions.

### Syntax

library liberrlog.a

```
#include <sys/errlog.h>
```

```
int errlog_set_direction(handle, direction)
errlog_handle_t handle;
int direction;
```

### Description

The `errlog_find_next` and `errlog_find_sequence` subroutines search the error log starting with the most recent log entry and going backward in time, by default. The `errlog_set_direction` subroutine is used to alter this direction.

### Parameters

The `handle` parameter must contain a handle returned by a previous `errlog_open` call.

The `direction` parameter must be `LE_FORWARD` or `LE_REVERSE`. `LE_REVERSE` is the default if the `errlog_set_direction` subroutine is not used.

### Return Values

Upon successful completion, the `errlog_set_direction` subroutine returns 0.

If a parameter is invalid, the `errlog_set_direction` subroutine returns `LE_ERR_INVARG`.

## Related Information

The `errlog_open`, `errlog_close`, `errlog_find_first`, `errlog_find_next`, `errlog_find_sequence`, `errlog_write`, and `errlog` subroutines.

---

## errlog\_write Subroutine

### Purpose

Changes the previously read error log entry.

### Syntax

library liberrlog.a

```
#include <sys/errlog.h>
```

```
int errlog_write(handle, entry)
errlog_handle_t handle;
errlog_entry_t *entry;
```

## Description

The **errlog\_write** subroutine is used to update the most recently read log entry. Neither the length nor the sequence number of the entry may be changed. The entry is simply updated in place.

If the **errlog\_write** subroutine is used in a multi-threaded application, the program should obtain a lock around the read/write pair to avoid conflict.

## Parameters

The handle parameter must contain a handle returned by a previous **errlog\_open** call.

The entry parameter must point to an entry returned by the previous error log find function.

## Return Values

Upon successful completion, the **errlog\_write** subroutine returns 0.

If a parameter is invalid, the **errlog\_write** subroutine returns **LE\_ERR\_INVARG**.

The **errlog\_write** subroutine may also return one of the following:

<b>LE_ERR_INVFILE</b>	The data on file is invalid.
<b>LE_ERR_IO</b>	An i/o error occurred.
<b>LE_ERR_NOWRITE</b>	The entry to be written didn't match the entry being updated.

## Related Information

The **errlog\_open**, **errlog\_close**, **errlog\_find\_first**, **errlog\_find\_next**, **errlog\_find\_sequence**, **errlog\_set\_direction**, and **errlog** subroutines.

The **/usr/include/sys/errlog.h** include file.

---

## exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine

### Purpose

Executes a file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
extern
char **environ;
```

```
int execl (
    Path,
```

```
Argument0 [, Argument1, ...], 0)
const char *Path, *Argument0, *Argument
1, ...;
```

```
int execl (
    Path,
    Argument0 [, Argument1, ...], 0,
```

```
EnvironmentPointer)
const
char *Path, *Argument0, *Argum
ent
1, ...;
char *const EnvironmentPointer[ ];
```

```
int execlp (
    File,
    Argument0 [, Argument1
, ...], 0)
const char *File, *Argument0, *Argument
1, ...;
```

```
int execv (
    Path,
    ArgumentV)
const char *Path;
char *const ArgumentV[ ];
```

```
int execve (
    Path,
    ArgumentV,
```

```
EnvironmentPointer)
const char *Path;
char
*const ArgumentV[ ], *EnvironmentPointer
[ ];
```

```
int execvp (
    File,
    ArgumentV)
const char *File;
char *const ArgumentV[ ];
```

```
int exect (
    Path,
    ArgumentV,
    EnvironmentPointer)
char *Path, *ArgumentV, *EnvironmentPointer [ ];
```

## Description

The **exec** subroutine, in all its forms, executes a new program in the calling process. The **exec** subroutine does not create a new process, but overlays the current program with a new one, which is called the *new-process image*. The new-process image file can be one of three file types:

- An executable binary file in XCOFF file format. .
- An executable text file that contains a shell procedure (only the **execlp** and **execvp** subroutines allow this type of new-process image file).
- A file that names an executable binary file or shell procedure to be run.

The new-process image inherits the following attributes from the calling process image: session membership, supplementary group IDs, process signal mask, and pending signals.

The last of the types mentioned is recognized by a header with the following syntax:

```
#! Path [String]
```

The **#!** is the file *magic number*, which identifies the file type. The path name of the file to be executed is specified by the *Path* parameter. The *String* parameter is an optional character string that contains no tab or space characters. If specified, this string is passed to the new process as an argument in front of the name of the new-process image file. The header must be terminated with a new-line character. When called, the new process passes the *Path* parameter as *ArgumentV[0]*. If a *String* parameter is specified in the new process image file, the **exec** subroutine sets *ArgumentV[0]* to the *String* and *Path* parameter values concatenated together. The rest of the arguments passed are the same as those passed to the **exec** subroutine.

The **exec** subroutine attempts to cancel outstanding asynchronous I/O requests by this process. If the asynchronous I/O requests cannot be canceled, the application is blocked until the requests have completed.

The **exec** subroutine is similar to the **load** (“load and loadAndInit Subroutines” on page 727) subroutine, except that the **exec** subroutine does not have an explicit library path parameter. Instead, the **exec** subroutine uses either the **LIBPATH** or **LD\_LIBRARY\_PATH** environment variable. The **LIBPATH** variable, when set, is used in favor of **LD\_LIBRARY\_PATH**; otherwise, **LD\_LIBRARY\_PATH** is used. These library path variables are ignored when the program that the **exec** subroutine is run on has more privilege than the calling program (for example, an **suid** program).

The **exec** subroutine is included for compatibility with older programs being traced with the **ptrace** command. The program being executed is forced into hardware single-step mode.

**Note:** **exec** is not supported in 64-bit mode.

**Note:** Currently, a Graphics Library program cannot be overlaid with another Graphics Library program. The overlaying program can be a nongraphics program. For additional information, see the **/usr/lpp/GL/README** file.

## Parameters

*Path*

Specifies a pointer to the path name of the new-process image file. If Network File System (NFS) is installed on your system, this path can cross into another node. Data is copied into local virtual memory before proceeding.

<i>File</i>	Specifies a pointer to the name of the new-process image file. Unless the <i>File</i> parameter is a full path name, the path prefix for the file is obtained by searching the directories named in the <b>PATH</b> environment variable. The initial environment is supplied by the shell. <b>Note:</b> The <b>execlp</b> subroutine and the <b>execvp</b> subroutine take <i>File</i> parameters, but the rest of the <b>exec</b> subroutines take <i>Path</i> parameters. (For information about the environment, see the <b>environment</b> miscellaneous facility and the <b>sh</b> command.)
<i>Argument0</i> [, <i>Argument1</i> , ...]	Point to null-terminated character strings. The strings constitute the argument list available to the new process. By convention, at least the <i>Argument0</i> parameter must be present, and it must point to a string that is the same as the <i>Path</i> parameter or its last component.
<i>ArgumentV</i>	Specifies an array of pointers to null-terminated character strings. These strings constitute the argument list available to the new process. By convention, the <i>ArgumentV</i> parameter must have at least one element, and it must point to a string that is the same as the <i>Path</i> parameter or its last component. The last element of the <i>ArgumentV</i> parameter is a null pointer.
<i>EnvironmentPointer</i>	An array of pointers to null-terminated character strings. These strings constitute the environment for the new process. The last element of the <i>EnvironmentPointer</i> parameter is a null pointer.

When a C program is run, it receives the following parameters:

```
main (ArgumentCount, ArgumentV, EnvironmentPointer)
int ArgumentCount;
char *ArgumentV[ ], *EnvironmentPointer[
];
```

In this example, the *ArgumentCount* parameter is the argument count, and the *ArgumentV* parameter is an array of character pointers to the arguments themselves. By convention, the value of the *ArgumentCount* parameter is at least 1, and the *ArgumentV*[0] parameter points to a string containing the name of the new-process image file.

The **main** routine of a C language program automatically begins with a runtime start-off routine. This routine sets the **environ** global variable so that it points to the environment array passed to the program in *EnvironmentPointer*. You can access this global variable by including the following declaration in your program:

```
extern char **environ;
```

The **execl**, **execv**, **execlp**, and **execvp** subroutines use the **environ** global variable to pass the calling process current environment to the new process.

File descriptors open in the calling process remain open, except for those whose **close-on-exec** flag is set. For those file descriptors that remain open, the file pointer is unchanged. (For information about file control, see the **fcntl.h** file.)

The state-of-conversion descriptors and message-catalog descriptors in the new process image are undefined. For the new process, an equivalent of the **setlocale** subroutine, specifying the **LC\_ALL** value for its category and the **"C"** value for its locale, is run at startup.

If the new program requires shared libraries, the **exec** subroutine finds, opens, and loads each of them into the new-process address space. The referenced counts for shared libraries in use by the issuer of the **exec** are decremented. Shared libraries are searched for in the directories listed in the **LIBPATH** environment variable. If any of these files is remote, the data is copied into local virtual memory.

The **exec** subroutines reset all caught signals to the default action. Signals that cause the default action continue to do so after the **exec** subroutines. Ignored signals remain ignored, the signal mask remains the same, and the signal stack state is reset. (For information about signals, see the **sigaction** subroutine.)

If the *SetUserID* mode bit of the new-process image file is set, the **exec** subroutine sets the effective user ID of the new process to the owner ID of the new-process image file. Similarly, if the *SetGroupID* mode bit of the new-process image file is set, the effective group ID of the new process is set to the group ID of the new-process image file. The real user ID and real group ID of the new process remain the same as those of the calling process. (For information about the *SetID* modes, see the **chmod** subroutine.)

At the end of the **exec** operation the saved user ID and saved group ID of the process are always set to the effective user ID and effective group ID, respectively, of the process.

When one or both of the set ID mode bits is set and the file to be executed is a remote file, the file user and group IDs go through outbound translation at the server. Then they are transmitted to the client node where they are translated according to the inbound translation table. These translated IDs become the user and group IDs of the new process.

**Note:** **setuid** and **setgid** bids on shell scripts do not affect user or group IDs of the process finally executed.

Profiling is disabled for the new process.

The new process inherits the following attributes from the calling process:

- Nice value (see the **getpriority** subroutine, **setpriority** subroutine, **nice** subroutine)
- Process ID
- Parent process ID
- Process group ID
- **semadj** values (see the **semop** subroutine)
- tty group ID (see the **exit**, **atexit**, or **\_exit** subroutine, **sigaction** subroutine)
- **trace** flag (see request 0 of the **ptrace** subroutine)
- Time left until an alarm clock signal (see the **incinterval** subroutine, **setitimer** subroutine, and **alarm** subroutine)
- Current directory
- Root directory
- File-mode creation mask (see the **umask** subroutine)
- File size limit (see the **ulimit** subroutine)
- Resource limits (see the **getrlimit** subroutine, **setrlimit** subroutine, and **vlimit** subroutine)
- **tms\_utime**, **tms\_stime**, **tms\_cutime**, and **tms\_ctime** fields of the **tms** structure (see the **times** subroutine)
- Login user ID

Upon successful completion, the **exec** subroutines mark for update the **st\_atime** field of the file.

## Examples

1. To run a command and pass it a parameter, enter:

```
exec1p("ls", "ls", "-al", 0);
```

The **exec1p** subroutine searches each of the directories listed in the **PATH** environment variable for the **ls** command, and then it overlays the current process image with this command. The **exec1p** subroutine is not returned, unless the **ls** command cannot be executed.



**Note:** This example does not run the shell command processor, so operations interpreted by the shell, such as using wildcard characters in file names, are not valid.

2. To run the shell to interpret a command, enter:

```
execl("/usr/bin/sh", "sh", "-c", "ls -l *.c",  
0);
```

This runs the **sh** command with the **-c** flag, which indicates that the following parameter is the command to be interpreted. This example uses the **execl** subroutine instead of the **execvp** subroutine because the full path name **/usr/bin/sh** is specified, making a path search unnecessary.

Running a shell command in a child process is generally more useful than simply using the **exec** subroutine, as shown in this example. The simplest way to do this is to use the **system** subroutine.

3. The following is an example of a new-process file that names a program to be run:

```
#!/usr/bin/awk -f  
{ for (i = NF; i > 0; --i) print $i }
```

If this file is named **reverse**, entering the following command on the command line:

```
reverse chapter1 chapter2
```

This command runs the following command:

```
/usr/bin/awk -f reverse chapter1 chapter2
```

**Note:** The **exec** subroutines use only the first line of the new-process image file and ignore the rest of it. Also, the **awk** command interprets the text that follows a # (pound sign) as a comment.

## Return Values

Upon successful completion, the **exec** subroutines do not return because the calling process image is overlaid by the new-process image. If the **exec** subroutines return to the calling process, the value of **-1** is returned and the **errno** global variable is set to identify the error.

## Error Codes

If the **exec** subroutine is unsuccessful, it returns one or more of the following error codes:

<b>EACCES</b>	The new-process image file is not an ordinary file.
<b>EACCES</b>	The mode of the new-process image file denies execution permission.
<b>ENOEXEC</b>	The <b>exec</b> subroutine is neither an <b>execvp</b> subroutine nor an <b>execvp</b> subroutine. The new-process image file has the appropriate access permission, but the magic number in its header is not valid.
<b>ENOEXEC</b>	The new-process image file has a valid magic number in its header, but the header is damaged or is incorrect for the machine on which the file is to be run.
<b>ETXTBSY</b>	The new-process image file is a pure procedure (shared text) file that is currently open for writing by some process.
<b>ENOMEM</b>	The new process requires more memory than is allowed by the system-imposed maximum, the <b>MAXMEM</b> compiler option.
<b>E2BIG</b>	The number of bytes in the new-process argument list is greater than the system-imposed limit. This limit is a system configurable value that can be set by superusers or system group users using SMIT. Refer to Kernel Tunable Parameters for details.
<b>EFAULT</b>	The <i>Path</i> , <i>ArgumentV</i> , or <i>EnvironmentPointer</i> parameter points outside of the process address space.
<b>EPERM</b>	The <i>SetUserID</i> or <i>SetGroupID</i> mode bit is set on the process image file. The translation tables at the server or client do not allow translation of this user or group ID.

If the **exec** subroutine is unsuccessful because of a condition requiring path name resolution, it returns one or more of the following error codes:

<b>EACCES</b>	Search permission is denied on a component of the path prefix. Access could be denied due to a secure mount.
<b>EFAULT</b>	The <i>Path</i> parameter points outside of the allocated address space of the process.
<b>EIO</b>	An input/output (I/O) error occurred during the operation.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of a path name exceeded 255 characters and the process has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine), or an entire path name exceeded 1023 characters.
<b>ENOENT</b>	A component of the path prefix does not exist.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOENT</b>	The path name is null.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>ESTALE</b>	The root or current directory of the process is located in a virtual file system that has been unmounted.

In addition, some errors can occur when using the new-process file after the old process image has been overwritten. These errors include problems in setting up new data and stack registers, problems in mapping a shared library, or problems in reading the new-process file. Because returning to the calling process is not possible, the system sends the **SIGKILL** signal to the process when one of these errors occurs.

If an error occurred while mapping a shared library, an error message describing the reason for error is written to standard error before the signal **SIGKILL** is sent to the process. If a shared library cannot be mapped, the subroutine returns one of the following error codes:

<b>ENOENT</b>	One or more components of the path name of the shared library file do not exist.
<b>ENOTDIR</b>	A component of the path prefix of the shared library file is not a directory.
<b>ENAMETOOLONG</b>	A component of a path name prefix of a shared library file exceeded 255 characters, or an entire path name exceeded 1023 characters.
<b>EACCES</b>	Search permission is denied for a directory listed in the path prefix of the shared library file.
<b>EACCES</b>	The shared library file mode denies execution permission.
<b>ENOEXEC</b>	The shared library file has the appropriate access permission, but a magic number in its header is not valid.
<b>ETXTBSY</b>	The shared library file is currently open for writing by some other process.
<b>ENOMEM</b>	The shared library requires more memory than is allowed by the system-imposed maximum.
<b>ESTALE</b>	The process root or current directory is located in a virtual file system that has been unmounted.
<b>EPROCLIM</b>	If WLM is running, the limit on the number of processes, threads, or logins in the class may have been met.

If NFS is installed on the system, the **exec** subroutine can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **alarm** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) or **incinterval** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) subroutine, **chmod** (“chmod or fchmod Subroutine” on page 148) or **fchmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fcntl** (“fcntl, dup, or dup2 Subroutine” on page 255) subroutine, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **getrusage** (“getrusage, getrusage64, times, or vtimes Subroutine” on page 428) or **times** (“getrusage, getrusage64,

times, or vtimes Subroutine” on page 428) subroutine, **nice** (“getpriority, setpriority, or nice Subroutine” on page 412) subroutine, **profil** (“profil Subroutine” on page 1176) subroutine, **ptrace** (“ptrace, ptracex, ptrace64 Subroutine” on page 1307) subroutine.

The “posix\_spawn or posix\_spawnnp Subroutine” on page 1149.

The **semop** subroutine, **settimer** (“gettimer, settimer, restimer, stime, or time Subroutine” on page 446) subroutine, **sigaction**, **signal**, or **sigvec** subroutine, **shmat** subroutine, **system** subroutine, **ulimit** subroutine, **umask** subroutine.

The **awk** command, **ksh** command, **sh** command.

The **environment** file.

The XCOFF object (**a.out**) file format.

The **varargs** macros.

Asynchronous I/O Overview in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

---

## exit, atexit, unatexit, \_exit, or \_Exit Subroutine

### Purpose

Terminates a process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
void exit ( Status)  
int Status;
```

```
void _exit ( Status)  
int Status;
```

```
void _Exit (Status)  
int Status;
```

```
#include <sys/limits.h>
```

```
int atexit ( Function)  
void (*Function) (void);
```

```
int unatexit (Function)  
void (*Function)(void);
```

### Description

The **exit** subroutine terminates the calling process after calling the standard I/O library **\_cleanup** function to flush any buffered output. Also, it calls any functions registered previously for the process by the **atexit** subroutine. The **atexit** subroutine registers functions called at normal process termination for cleanup processing. Normal termination occurs as a result of either a call to the **exit** subroutine or a **return** statement in the **main** function.

Each function a call to the **atexit** subroutine registers must return. This action ensures that all registered functions are called.

Finally, the **exit** subroutine calls the **\_exit** subroutine, which completes process termination and does not return. The **\_exit** subroutine terminates the calling process and causes the following to occur:

The **\_Exit** subroutine is functionally equivalent to the **\_exit** subroutine. The **\_Exit** subroutine does not call functions registered with **atexit** or any registered signal handlers. The way the subroutine is implemented determines whether open streams are flushed or closed, and whether temporary files are removed. The calling process is terminated with the consequences described below.

- All of the file descriptors, directory streams, conversion descriptors, and message catalog descriptors open in the calling process are closed.
- If the parent process of the calling process is executing a **wait** or **waitpid**, and has not set its **SA\_NOCLDWAIT** flag nor set **SIGCHLD** to **SIG\_IGN**, it is notified of the calling process' termination and the low order eight bits (that is, bits 0377) of *status* are made available to it. If the parent is not waiting, the child's status is made available to it when the parent subsequently executes **wait** or **waitpid**.
- If the parent process of the calling process is not executing a **wait** or **waitpid**, and has neither set its **SA\_NOCLDWAIT** flag nor set **SIGCHLD** to **SIG\_IGN**, the calling process is transformed into a zombie process. A zombie process is an inactive process that is deleted at some later time when its parent process executes **wait** or **waitpid**.
- Termination of a process does not directly terminate its children. The sending of a **SIGHUP** signal indirectly terminates children in some circumstances. This can be accomplished in one of two ways. If the implementation supports the **SIGCHLD** signal, a **SIGCHLD** is sent to the parent process. If the parent process has set its **SA\_NOCLDWAIT** flag, or set **SIGCHLD** to **SIG\_IGN**, the status is discarded, and the lifetime of the calling process ends immediately. If **SA\_NOCLDWAIT** is set, it is implementation defined whether a **SIGCHLD** signal is sent to the parent process.
- The parent process ID of all of the calling process' existing child processes and zombie processes are set to the process ID of an implementation defined system process.
- Each attached shared memory segment is detached and the value of *shm\_nattch* (see **shmget**) in the data structure associated with its shared memory ID is decremented by 1.
- For each semaphore for which the calling process has set a *semadj* value (see **semop**), that value is added to the *semval* of the specified semaphore.
- If the process is a controlling process, the **SIGHUP** signal is sent to each process in the foreground process group of the controlling terminal belonging to the calling process.
- If the process is a controlling process, the controlling terminal associated with the session is disassociated from the session, allowing it to be acquired by a new controlling process.
- If the exit of the process causes a process group to become orphaned, and if any member of the newly orphaned process group is stopped, a **SIGHUP** signal followed by a **SIGCONT** signal is sent to each process in the newly orphaned process group.
- All open named semaphores in the calling process are closed as if by appropriate calls to **sem\_close**.
- Memory mappings that were created in the process are unmapped before the process is destroyed.
- Any blocks of typed memory that were mapped in the calling process are unmapped, as if the **munmap** subroutine was implicitly called to unmap them.
- All open message queue descriptors in the calling process are closed.
- Any outstanding cancelable asynchronous I/O operations may be canceled. Those asynchronous I/O operations that are not canceled complete as if the **\_Exit** subroutine had not yet occurred, but any associated signal notifications are suppressed.

The **\_Exit** subroutine may block awaiting such I/O completion. The implementation defines whether any I/O is canceled, and which I/O may be canceled upon **\_Exit**.

- Threads terminated by a call to **\_Exit** do not invoke their cancelation cleanup handlers or per thread data destructors.

- If the calling process is a trace controller process, any trace streams that were created by the calling process are shut down.

The **unatexit()** subroutine is used to unregister functions that were previously registered by the **atexit()** subroutine. If the referenced function is found, it is removed from the list of functions that are called at normal program termination.

## Parameters

<i>Status</i>	Indicates the status of the process. May be set to 0, EXIT_SUCCESS, EXIT_FAILURE, or any other value, though only the least significant 8 bits are available to a waiting parent process.
<i>Function</i>	Specifies a function to be called at normal process termination for cleanup processing. You may specify a number of functions to the limit set by the <b>ATEXIT_MAX</b> function, which is defined in the <b>sys/limits.h</b> file. A pushdown stack of functions is kept so that the last function registered is the first function called.

## Return Values

Upon successful completion, the **atexit** subroutine returns a value of 0. Otherwise, a nonzero value is returned. The **exit** and **\_exit** subroutines do not return a value.

The **unatexit()** subroutine returns a value of 0 if the function referenced by *Function* is found and removed from the **atexit** list. Otherwise, a non-zero value is returned.

## Related Information

“acct Subroutine” on page 7, “lockfx, lockf, flock, or lockf64 Subroutine” on page 740, “lockfx, lockf, flock, or lockf64 Subroutine” on page 740, “lockfx, lockf, flock, or lockf64 Subroutine” on page 740, and “getrusage, getrusage64, times, or vtimes Subroutine” on page 428.

**longjmp Subroutine, semop Subroutine, shmget Subroutine, sigaction, sigvec, or signal Subroutine, and wait, waitpid, or wait3 Subroutine** in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Asynchronous I/O Overview in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

**unistd.h** in *AIX 5L Version 5.3 Files Reference*.

## exp, expf, or expl Subroutine

### Purpose

Computes exponential, logarithm, and power functions.

### Libraries

IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
```

```
double exp ( x )
double x;
```

```
float expf (x)
float x;
long double expl (x)
long double x;
```

## Description

These subroutines are used to compute exponential, logarithm, and power functions.

The **exp**, **expf**, and **expl** subroutines returns  $\exp(x)$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies some double-precision floating-point value.  
*y* Specifies some double-precision floating-point value.

## Return Values

Upon successful completion, the **exp**, **expf**, and **expl** subroutines return the exponential value of *x*.

If the correct value would cause overflow, a range error occurs and the **exp**, **expf**, and **expl** subroutine returns the value of the macro **HUGE\_VAL**, **HUGE\_VALF** and **HUGE\_VALL**, respectively.

If the correct value would cause underflow, and is not representable, a range error may occur, and either 0.0 (if supported), or an implementation-defined value is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ , 1 is returned.

If *x* is  $-\text{Inf}$ ,  $+\text{0}$  is returned.

If *x* is  $+\text{Inf}$ , *x* is returned.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

## Error Codes

When using the **libm.a** library:

**exp** If the correct value would overflow, the **exp** subroutine returns a **HUGE\_VAL** value and the **errno** global variable is set to a **ERANGE** value.

When using **libmsaa.a(-lmsaa)**:

**exp** If the correct value would overflow, the **exp** subroutine returns a **HUGE\_VAL** value. If the correct value would underflow, the **exp** subroutine returns 0. In both cases **errno** is set to **ERANGE**.

**expl** If the correct value would overflow, the **expl** subroutine returns a **HUGE\_VAL** value. If the correct value would underflow, the **expl** subroutine returns 0. In both cases **errno** is set to **ERANGE**.

**expl** If the correct value overflows, the **expl** subroutine returns a **HUGE\_VAL** value and **errno** is set to **ERANGE**.

These error-handling procedures may be changed with the **matherr** subroutine when using the **libmsaa.a** library.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

The **matherr** (“matherr Subroutine” on page 789) subroutine, **sinh**, **cosh**, or **tanh** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## exp2, exp2f, or exp2l Subroutine

### Purpose

Computes the base 2 exponential.

### Syntax

```
#include <math.h>
```

```
double exp2 (x)  
double x;
```

```
float exp2f (x)  
float x;
```

```
long double exp2l (x)  
long double x;
```

### Description

The **exp2**, **exp2f**, and **exp2l** subroutines compute the base 2 exponential of the *x* parameter.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept** (**FE\_ALL\_EXCEPT**) before calling these subroutines. On return, if **errno** is nonzero or **fetestexcept** (**FE\_INVALID** | **FE\_DIVBYZERO** | **FE\_OVERFLOW** | **FE\_UNDERFLOW**) is nonzero, an error has occurred.

### Parameters

*x* Specifies the base 2 exponential to be computed.

### Return Values

Upon successful completion, the **exp2**, **exp2f**, or **exp2l** subroutine returns  $2^x$ .

If the correct value causes overflow, a range error occurs and the **exp2**, **exp2f**, and **exp2l** subroutines return the value of the macro (**HUGE\_VAL**, **HUGE\_VALF**, and **HUGE\_VALL**, respectively).

If the correct value causes underflow and is not representable, a range error occurs, and 0.0 is returned.

If  $x$  is NaN, NaN is returned.

If  $x$  is  $\pm 0$ , 1 is returned.

If  $x$  is  $-\text{Inf}$ , 0 is returned.

If  $x$  is  $+\text{Inf}$ ,  $x$  is returned.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

## Related Information

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## expm1, expm1f, or expm1l Subroutine

### Purpose

Computes exponential functions.

### Syntax

```
#include <math.h>
```

```
float expm1f (x)
float x;
```

```
long double expm1l (x)
long double x;
```

```
double expm1 (x)
double x;
```

### Description

The **expm1f**, **expm1l**, and **expm1** subroutines compute  $e^x - 1.0$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

### Return Values

Upon successful completion, the **expm1f**, **expm1l**, and **expm1** subroutines return  $e^x - 1.0$ .

If the correct value would cause overflow, a range error occurs and the **expm1f**, **expm1l**, and **expm1** subroutines return the value of the macro **HUGE\_VALF**, **HUGE\_VALL**, and **HUGE\_VAL**, respectively.



If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$ ,  $\pm 0$  is returned.

If  $x$  is  $-\text{Inf}$ ,  $-1$  is returned.

If  $x$  is  $+\text{Inf}$ ,  $x$  is returned.

If  $x$  is subnormal, a range error may occur and  $x$  is returned.

## Related Information

“exp, expf, or expl Subroutine” on page 245, “feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, “ilogbf, ilogbl, or ilogb Subroutine” on page 534, and “log, logf, or logl Subroutine” on page 749.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## **fabsf, fabsf, fabs, fabsd32, fabsd64, and fabsd128 Subroutines**

### Purpose

Determines the absolute value.

### Syntax

```
#include <math.h>
```

```
float fabsf (x)  
float x;
```

```
long double fabsl (x)  
long double x;
```

```
double fabs (x)  
double x;
```

```
_Decimal32 fabsd32 (x)  
_Decimal32 x;
```

```
_Decimal64 fabsd64 (x)  
_Decimal64 x;
```

```
_Decimal128 fabsd128 (x)  
_Decimal128 x;
```

### Description

The **fabsf**, **fabsl**, **fabs**, **fabsd32**, **fabsd64**, and **fabsd128** subroutines compute the absolute value of the  $x$  parameter,  $|x|$ .

### Parameters

$x$  Specifies the value to be computed.

### Return Values

Upon successful completion, the **fabsf**, **fabsl**, **fabs**, **fabsd32**, **fabsd64**, and **fabsd128** subroutines return the absolute value of  $x$ .

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$ ,  $+0$  is returned.

If  $x$  is  $\pm \text{Inf}$ ,  $+\text{Inf}$  is returned.

## Related Information

The “class, `_class`, finite, isnan, or unordered Subroutines” on page 167.

`math.h` in *AIX 5L Version 5.3 Files Reference*.

---

## fattach Subroutine

### Purpose

Attaches a STREAMS-based file descriptor to a file.

### Library

Standard C Library (`libc.a`)

### Syntax

```
#include <stropts.h>
int fattach(int fildev, const char *path);
```

### Description

The `fattach` subroutine attaches a STREAMS-based file descriptor to a file, effectively associating a pathname with *fildev*. The *fildev* argument must be a valid open file descriptor associated with a STREAMS file. The *path* argument points to a pathname of an existing file. The process must have appropriate privileges, or must be the owner of the file named by *path* and have write permission. A successful call to `fattach` subroutine causes all pathnames that name the file named by *path* to name the STREAMS file associated with *fildev*, until the STREAMS file is detached from the file. A STREAMS file can be attached to more than one file and can have several pathnames associated with it.

The attributes of the named STREAMS file are initialized as follows: the permissions, user ID, group ID, and times are set to those of the file named by *path*, the number of links is set to 1, and the size and device identifier are set to those of the STREAMS file associated with *fildev*. If any attributes of the named STREAMS file are subsequently changed (for example, by `chmod` subroutine), neither the attributes of the underlying file nor the attributes of the STREAMS file to which *fildev* refers are affected.

File descriptors referring to the underlying file, opened prior to an `fattach` subroutine, continue to refer to the underlying file.

### Parameters

<i>fildev</i>	A file descriptor identifying an open STREAMS-based object.
<i>path</i>	An existing pathname which will be associated with <i>fildev</i> .

### Return Value

0	Successful completion.
-1	Not successful and <i>errno</i> set to one of the following.

## Errno Value

<b>EACCES</b>	Search permission is denied for a component of the path prefix, or the process is the owner of <i>path</i> but does not have write permission on the file named by <i>path</i> .
<b>EBADF</b>	The file referred to by <i>filides</i> is not an open file descriptor.
<b>ENOENT</b>	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>EPERM</b>	The effective user ID of the process is not the owner of the file named by <i>path</i> and the process does not have appropriate privilege.
<b>EBUSY</b>	The file named by <i>path</i> is currently a mount point or has a STREAMS file attached to it.
<b>ENAMETOOLONG</b>	The size of <i>path</i> exceeds <b>{PATH_MAX}</b> , or a component of <i>path</i> is longer than <b>{NAME_MAX}</b> .
<b>ELOOP</b>	Too many symbolic links were encountered in resolving <i>path</i> .
<b>EINVAL</b>	The <i>filides</i> argument does not refer to a STREAMS file.
<b>ENOMEM</b>	Insufficient storage space is available.

## Related Specifics

The **fdetach** (“fdetach Subroutine” on page 261) subroutine, **isastream** subroutine.

---

## fchdir Subroutine

### Purpose

Directory pointed to by the file descriptor becomes the current working directory.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int fchdir (int Fildes)
```

### Description

The **fchdir** subroutine causes the directory specified by the *Fildes* parameter to become the current working directory.

### Parameter

*Fildes*            A file descriptor identifying an open directory obtained from a call to the **open** subroutine.

### Return Values

0            Successful completion  
-1           Not successful and **errno** set to one of the following.

### Error Codes

<b>EACCES</b>	Search access is denied.
<b>EBADF</b>	The file referred to by <i>Fildes</i> is not an open file descriptor.
<b>ENOTDIR</b>	The open file descriptor does not refer to a directory.

## Related Information

The **chdir** (“chdir Subroutine” on page 147) subroutine, **chroot** (“chroot Subroutine” on page 160) subroutine, **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine.

---

## fclear or fclear64 Subroutine

### Purpose

Makes a hole in a file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
off_t fclear ( FileDescriptor, NumberOfBytes )  
int FileDescriptor;  
off_t NumberOfBytes;
```

```
off64_t fclear64 ( FileDescriptor, NumberOfBytes )  
int FileDescriptor;  
off64_t NumberOfBytes;
```

### Description

The **fclear** and **fclear64** subroutines zero the number of bytes specified by the *NumberOfBytes* parameter starting at the current file pointer for the file specified in the *FileDescriptor* parameter. If Network File System (NFS) is installed on your system, this file can reside on another node.

The **fclear** subroutine can only clear up to **OFF\_MAX** bytes of the file while **fclear64** can clear up to the maximum file size.

The **fclear** and **fclear64** subroutines cannot be applied to a file that a process has opened with the **O\_DEFER** mode.

Successful completion of the **fclear** and **fclear64** subroutines clear the SetUserID bit (**S\_ISUID**) of the file if any of the following are true:

- The calling process does not have root user authority.
- The effective user ID of the calling process does not match the user ID of the file.
- The file is executable by the group (**S\_IXGRP**) or others (**S\_IXOTH**).

This subroutine also clears the SetGroupID bit (**S\_ISGID**) if:

- The file does not match the effective group ID or one of the supplementary group IDs of the process,  
OR
- The file is executable by the owner (**S\_IXUSR**) or others (**S\_IXOTH**).

**Note:** Clearing of the SetUserID and SetGroupID bits can occur even if the subroutine fails because the data in the file was modified before the error was detected.

In the large file enabled programming environment, **fclear** is redefined to be **fclear64**.

## Parameters

<i>FileDescriptor</i>	Indicates the file specified by the <i>FileDescriptor</i> parameter must be open for writing. The <i>FileDescriptor</i> is a small positive integer used instead of the file name to identify a file. This function differs from the logically equivalent write operation in that it returns full blocks of binary zeros to the file system, constructing holes in the file.
<i>NumberOfBytes</i>	Indicates the number of bytes that the seek pointer is advanced. If you use the <b>fclear</b> and <b>fclear64</b> subroutines past the end of a file, the rest of the file is cleared and the seek pointer is advanced by <i>NumberOfBytes</i> . The file size is updated to include this new hole, which leaves the current file position at the byte immediately beyond the new end-of-file pointer.

## Return Values

Upon successful completion, a value of *NumberOfBytes* is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **fclear** and **fclear64** subroutines fail if one or more of the following are true:

<b>EIO</b>	I/O error.
<b>EBADF</b>	The <i>FileDescriptor</i> value is not a valid file descriptor open for writing.
<b>EINVAL</b>	The file is not a regular file.
<b>EMFILE</b>	The file is mapped <b>O_DEFER</b> by one or more processes.
<b>EAGAIN</b>	The write operation in the <b>fclear</b> subroutine failed due to an enforced write lock on the file.
<b>EFBIG</b>	The current offset plus <i>NumberOfBytes</i> exceeds the offset maximum established in the open file description associated with <i>FileDescriptor</i> .
<b>EFBIG</b>	An attempt was made to write a file that exceeds the process' file size limit or the maximum file size. If the user has set the environment variable <b>XPG_SUS_ENV=ON</b> prior to execution of the process, then the <b>SIGXFSZ</b> signal is posted to the process when exceeding the process' file size limit.

If NFS is installed on the system the **fclear** and **fclear64** subroutines can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **open**, **openx**, or **creat** ("open, openx, open64, creat, or creat64 Subroutine" on page 935) subroutine, **truncate** or **ftruncate** subroutines.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fclose or fflush Subroutine

### Purpose

Closes or flushes a stream.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdio.h>
```

```
int fclose ( Stream)  
FILE *Stream;
```

```
int fflush ( Stream)  
FILE *Stream;
```

## Description

The **fclose** subroutine writes buffered data to the stream specified by the *Stream* parameter, and then closes the stream. The **fclose** subroutine is automatically called for all open files when the **exit** subroutine is invoked.

The **fflush** subroutine writes any buffered data for the stream specified by the *Stream* parameter and leaves the stream open. The **fflush** subroutine marks the *st\_ctime* and *st\_mtime* fields of the underlying file for update.

If the *Stream* parameter is a null pointer, the **fflush** subroutine performs this flushing action on all streams for which the behavior is defined.

## Parameters

*Stream*                      Specifies the output stream.

## Return Values

Upon successful completion, the **fclose** and **fflush** subroutines return a value of 0. Otherwise, a value of EOF is returned.

## Error Codes

If the **fclose** and **fflush** subroutines are unsuccessful, the following errors are returned through the **errno** global variable:

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> or <b>O_NDELAY</b> flag is set for the file descriptor underlying the <i>Stream</i> parameter and the process would be delayed in the write operation.
<b>EBADF</b>	The file descriptor underlying <i>Stream</i> is not valid.
<b>EFBIG</b>	An attempt was made to write a file that exceeds the process' file size limit or the maximum file size. See the <b>ulimit</b> subroutine.
<b>EFBIG</b>	The file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.
<b>EINTR</b>	The <b>fflush</b> subroutine was interrupted by a signal.
<b>EIO</b>	The process is a member of a background process group attempting to write to its controlling terminal, the <b>TOSTOP</b> signal is set, the process is neither ignoring nor blocking the <b>SIGTTOU</b> signal and the process group of the process is orphaned. This error may also be returned under implementation-dependent conditions.
<b>ENOSPC</b>	No free space remained on the device containing the file.
<b>EPIPE</b>	An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <b>SIGPIPE</b> signal is sent to the process.
<b>ENXIO</b>	A request was made of a non-existent device, or the request was outside the capabilities of the device

## Related Information

The **close** (“close Subroutine” on page 175) subroutine, **exit**, **atexit**, or **\_exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **setbuf**, **setvbuf**, **setbuffer**, or **setlinebuf** subroutine.

Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fcntl, dup, or dup2 Subroutine

### Purpose

Controls open file descriptors.

### Library

Standard C Library (**libc.a**)

Berkeley compatibility library (**libbsd.a**) (for the **fcntl** subroutine)

### Syntax

```
#include <fcntl.h>
```

```
int fcntl ( FileDescriptor, Command, Argument)  
int FileDescriptor, Command, Argument;
```

```
#include <unistd.h>
```

```
int dup2( Old, New)  
int Old, New;
```

```
int dup( FileDescriptor)  
int FileDescriptor;
```

### Description

The **fcntl** subroutine performs controlling operations on the open file specified by the *FileDescriptor* parameter. If Network File System (NFS) is installed on your system, the open file can reside on another node. The **fcntl** subroutine is used to:

- Duplicate open file descriptors.
- Set and get the file-descriptor flags.
- Set and get the file-status flags.
- Manage record locks.
- Manage asynchronous I/O ownership.
- Close multiple files.

The **fcntl** subroutine can provide the same functions as the **dup** and **dup2** subroutines.

If *FileDescriptor* refers to a terminal device or socket, then asynchronous I/O facilities can be used. These facilities are normally enabled by using the **ioctl** subroutine with the **FIOASYNC**, **FIOSETOWN**, and **FIOGETOWN** commands. However, a BSD-compatible mechanism is also available if the application is linked with the **libbsd.a** library.

When the *FileDescriptor* parameter refers to a shared memory object, the **fcntl** subroutine manages only the **F\_DUPFD**, **F\_DUP2FD**, **F\_GETFD**, **F\_SETFD**, **F\_GETFL**, and **F\_CLOSEM** commands.

When using the **libbsd.a** library, asynchronous I/O is enabled by using the **F\_SETFL** command with the **FASYNC** flag set in the *Argument* parameter. The **F\_GETOWN** and **F\_SETOWN** commands get the current asynchronous I/O owner and set the asynchronous I/O owner. However, these commands are valid only when the file descriptor refers to a terminal device or a socket.

All applications containing the **fcntl** subroutine must be compiled with **\_BSD** set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## General Record Locking Information

A lock is either an *enforced* or *advisory lock* and either a *read* or a *write lock*.

**Attention:** Buffered I/O does not work properly when used with file locking. Do not use the standard I/O package routines on files that are going to be locked.

For a lock to be an enforced lock, the Enforced Locking attribute of the file must be set; for example, the **S\_ENFMT** bit must be set, but the **S\_IXGRP**, **S\_IXUSR**, and **S\_IXOTH** bits must be clear. Otherwise, the lock is an advisory lock. A given file can have advisory or enforced locks, but not both. The description of the **sys/mode.h** file includes a description of file attributes.

When a process holds an enforced lock on a section of a file, no other process can access that section of the file with the **read** or **write** subroutine. In addition, the **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) and **ftruncate** subroutines cannot truncate the locked section of the file, and the **fclear** (“fclear or fclear64 Subroutine” on page 252) subroutine cannot modify the locked section of the file. If another process attempts to read or modify the locked section of the file, the process either sleeps until the section is unlocked or returns with an error indication.

When a process holds an advisory lock on a section of a file, no other process can lock that section of the file (or an overlapping section) with the **fcntl** subroutine. (No other subroutines are affected.) As a result, processes must voluntarily call the **fcntl** subroutine in order to make advisory locks effective.

When a process holds a read lock on a section of a file, other processes can also set read locks on that section or on subsets of it. Read locks are also called *shared* locks.

A read lock prevents any other process from setting a write lock on any part of the protected area. If the read lock is also an enforced lock, no other process can modify the protected area.

The file descriptor on which a read lock is being placed must have been opened with read access.

When a process holds a write lock on a section of a file, no other process can set a read lock or a write lock on that section. Write locks are also called *exclusive* locks. Only one write lock and no read locks can exist for a specific section of a file at any time.

If the lock is also an enforced lock, no other process can read or modify the protected area.

The following general rules about file locking apply:

- Changing or unlocking part of a file in the middle of a locked section leaves two smaller sections locked at each end of the originally locked section.
- If the calling process holds a lock on a file, that lock can be replaced by later calls to the **fcntl** subroutine.
- All locks associated with a file for a given process are removed when the process closes *any* file descriptor for that file.



- Locks are not inherited by a child process after a **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine is run.

**Note:** Deadlocks due to file locks in a distributed system are not always detected. When such deadlocks can possibly occur, the programs requesting the locks should set time-out timers.

Locks can start and extend beyond the current end of a file but cannot be negative relative to the beginning of the file. A lock can be set to extend to the end of the file by setting the `l_len` field to 0. If such a lock also has the `l_start` and `l_whence` fields set to 0, the whole file is locked. The `l_len`, `l_start`, and `l_whence` locking fields are part of the **flock** structure.

**Note:** The following description applies to AIX 4.3 and later releases.

When an application locks a region of a file using the 32 bit locking interface (`F_SETLK`), and the last byte of the lock range includes `MAX_OFF` (2 Gb - 1), then the lock range for the unlock request will be extended to include `MAX_END` ( $2^{63} - 1$ ).

## Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor obtained from a successful call to the <b>open</b> subroutine, <b>fcntl</b> subroutine, <b>pipe</b> subroutine, or <b>shm_open</b> subroutine. File descriptors are small positive integers used (instead of file names) to identify files or a shared memory object.
<i>Argument</i>	Specifies a variable whose value sets the function specified by the <i>Command</i> parameter. When dealing with file locks, the <i>Argument</i> parameter must be a pointer to the <b>FLOCK</b> structure.
<i>Command</i>	Specifies the operation performed by the <b>fcntl</b> subroutine. The <b>fcntl</b> subroutine can duplicate open file descriptors, set file-descriptor flags, set file descriptor locks, set process IDs, and close open file descriptors.

## Duplicating File Descriptors

<b>F_DUPFD</b>	Returns a new file descriptor as follows: <ul style="list-style-type: none"> <li>• Lowest-numbered available file descriptor greater than or equal to the <i>Argument</i> parameter</li> <li>• Same object references as the original file</li> <li>• Same file pointer as the original file (that is, both file descriptors share one file pointer if the object is a file)</li> <li>• Same access mode (read, write, or read-write)</li> <li>• Same file status flags (That is, both file descriptors share the same file status flags.)</li> <li>• The <b>close-on-exec</b> flag (<b>FD_CLOEXEC</b> bit) associated with the new file descriptor is cleared</li> </ul>
----------------	---

## Setting File-Descriptor Flags

<b>F_GETFD</b>	Gets the <b>close-on-exec</b> flag ( <b>FD_CLOEXEC</b> bit) that is associated with the file descriptor specified by the <i>FileDescriptor</i> parameter. The <i>Argument</i> parameter is ignored. File descriptor flags are associated with a single file descriptor, and do not affect others associated with the same file.
<b>F_SETFD</b>	Assigns the value of the <i>Argument</i> parameter to the <b>close-on-exec</b> flag ( <b>FD_CLOEXEC</b> bit) that is associated with the <i>FileDescriptor</i> parameter. If the <b>FD_CLOEXEC</b> flag value is 0, the file remains open across any calls to <b>exec</b> subroutines; otherwise, the file will close upon the successful execution of an <b>exec</b> subroutine.

**F\_GETFL** Gets the file-status flags and file-access modes for the open file description associated with the file descriptor specified by the *FileDescriptor* parameter. The open file description is set at the time the file is opened and applies only to those file descriptors associated with that particular call to the file. This open file descriptor does not affect other file descriptors that refer to the same file with different open file descriptions.

The file-status flags have the following values:

**O\_APPEND**

Set append mode.

**O\_NONBLOCK**

No delay.

The file-access modes have the following values:

**O\_RDONLY**

Open for reading only.

**O\_RDWR**

Open for reading and writing.

**O\_WRONLY**

Open for writing only.

The file access flags can be extracted from the return value using the **O\_ACCMODE** mask, which is defined in the **fcntl.h** file.

**F\_SETFL** Sets the file status flags from the corresponding bits specified by the *Argument* parameter. The file-status flags are set for the open file description associated with the file descriptor specified by the *FileDescriptor* parameter. The following flags may be set:

- **O\_APPEND** or **FAPPEND**
- **O\_NDELAY** or **FNDELAY**
- **O\_NONBLOCK** or **FNONBLOCK**
- **O\_SYNC** or **FSYNC**
- **FASYNC**

The **O\_NDELAY** and **O\_NONBLOCK** flags affect only operations against file descriptors derived from the same **open** subroutine. In BSD, these operations apply to all file descriptors that refer to the object.

## Setting File Locks

**F\_GETLK** Gets information on the first lock that blocks the lock described in the **flock** structure. The *Argument* parameter should be a pointer to a type **struct flock**, as defined in the **flock.h** file. The information retrieved by the **fcntl** subroutine overwrites the information in the **struct flock** pointed to by the *Argument* parameter. If no lock is found that would prevent this lock from being created, the structure is left unchanged, except for lock type (*l\_type*) which is set to **F\_UNLCK**.

**F\_SETLK** Sets or clears a file-segment lock according to the lock description pointed to by the *Argument* parameter. The *Argument* parameter should be a pointer to a type **struct flock**, which is defined in the **flock.h** file. The **F\_SETLK** option is used to establish read (or shared) locks (**F\_RDLCK**), or write (or exclusive) locks (**F\_WRLCK**), as well as to remove either type of lock (**F\_UNLCK**). The lock types are defined by the **fcntl.h** file. If a shared or exclusive lock cannot be set, the **fcntl** subroutine returns immediately.

**F\_SETLKW** Performs the same function as the **F\_SETLK** option unless a read or write lock is blocked by existing locks, in which case the process sleeps until the section of the file is free to be locked. If a signal that is to be caught is received while the **fcntl** subroutine is waiting for a region, the **fcntl** subroutine is interrupted, returns a -1, sets the **errno** global variable to **EINTR**. The lock operation is not done.

- F\_GETLK64** Gets information on the first lock that blocks the lock described in the **flock64** structure. The *Argument* parameter should be a pointer to an object of the type **struct flock64**, as defined in the **flock.h** file. The information retrieved by the **fcntl** subroutine overwrites the information in the **struct flock64** pointed to by the *Argument* parameter. If no lock is found that would prevent this lock from being created, the structure is left unchanged, except for lock type (*l\_type*) which is set to **F\_UNLCK**.
- F\_SETLK64** Sets or clears a file-segment lock according to the lock description pointed to by the *Argument* parameter. The *Argument* parameter should be a pointer to a type **struct flock64**, which is defined in the **flock.h** file. The **F\_SETLK** option is used to establish read (or shared) locks (**F\_RDLCK**), or write (or exclusive) locks (**F\_WRLCK**), as well as to remove either type of lock (**F\_UNLCK**). The lock types are defined by the **fcntl.h** file. If a shared or exclusive lock cannot be set, the **fcntl** subroutine returns immediately.
- F\_SETLKW64** Performs the same function as the **F\_SETLK** option unless a read or write lock is blocked by existing locks, in which case the process sleeps until the section of the file is free to be locked. If a signal that is to be caught is received while the **fcntl** subroutine is waiting for a region, the **fcntl** subroutine is interrupted, returns a -1, sets the **errno** global variable to **EINTR**. The lock operation is not done.

## Setting Process ID

- F\_GETOWN** Gets the process ID or process group currently receiving **SIGIO** and **SIGURG** signals. Process groups are returned as negative values.
- F\_SETOWN** Sets the process or process group to receive **SIGIO** and **SIGURG** signals. Process groups are specified by supplying a negative *Argument* value. Otherwise, the *Argument* parameter is interpreted as a process ID.

## Closing File Descriptors

- F\_CLOSEM** Closes all file descriptors from *FileDescriptor* up to the number specified by the **OPEN\_MAX** value.
- Old* Specifies an open file descriptor.
- New* Specifies an open file descriptor that is returned by the **dup2** subroutine.

## Compatibility Interfaces

### The lockfx Subroutine

The **fcntl** subroutine functions similar to the **lockfx** subroutine, when the *Command* parameter is **F\_SETLK**, **F\_SETLKW**, or **F\_GETLK**, and when used in the following way:

**fcntl** (*FileDescriptor*, *Command*, *Argument*)

is equivalent to:

**lockfx** (*FileDescriptor*, *Command*, *Argument*)

### The dup and dup2 Subroutines

The **fcntl** subroutine functions similar to the **dup** and **dup2** subroutines, when used in the following way:

**dup** (*FileDescriptor*)

is equivalent to:

**fcntl** (*FileDescriptor*, **F\_DUPFD**, 0)

**dup2** (*Old*, *New*)

is equivalent to:

```
close (New);
fcntl(Old, F_DUPFD, New)
```

The **dup** and **dup2** subroutines differ from the **fcntl** subroutine in the following ways:

- If the file descriptor specified by the *New* parameter is greater than or equal to **OPEN\_MAX**, the **dup2** subroutine returns a -1 and sets the **errno** variable to **EBADF**.
- If the file descriptor specified by the *Old* parameter is valid and equal to the file descriptor specified by the *New* parameter, the **dup2** subroutine will return the file descriptor specified by the *New* parameter, without closing it.
- If the file descriptor specified by the *Old* parameter is not valid, the **dup2** subroutine will be unsuccessful and will not close the file descriptor specified by the *New* parameter.
- The value returned by the **dup** and **dup2** subroutines is equal to the *New* parameter upon successful completion; otherwise, the return value is -1.

## Return Values

Upon successful completion, the value returned depends on the value of the *Command* parameter, as follows:

Command	Return Value
<b>F_DUPFD</b>	A new file descriptor
<b>F_GETFD</b>	The value of the flag (only the <b>FD_CLOEXEC</b> bit is defined)
<b>F_SETFD</b>	A value other than -1
<b>F_GETFL</b>	The value of file flags
<b>F_SETFL</b>	A value other than -1
<b>F_GETOWN</b>	The value of descriptor owner
<b>F_SETOWN</b>	A value other than -1
<b>F_GETLK</b>	A value other than -1
<b>F_SETLK</b>	A value other than -1
<b>F_SETLKW</b>	A value other than -1
<b>F_CLOSEM</b>	A value other than -1.

If the **fcntl** subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **fcntl** subroutine is unsuccessful if one or more of the following are true:

<b>EACCES</b>	The <i>Command</i> argument is <b>F_SETLK</b> ; the type of lock is a shared or exclusive lock and the segment of a file to be locked is already exclusively-locked by another process, or the type is an exclusive lock and some portion of the segment of a file to be locked is already shared-locked or exclusive-locked by another process.
<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid open file descriptor.
<b>EDEADLK</b>	The <i>Command</i> argument is <b>F_SETLKW</b> ; the lock is blocked by some lock from another process and putting the calling process to sleep, waiting for that lock to become free would cause a deadlock.
<b>ENOTTY</b>	The file descriptor does not refer to a terminal device or socket.
<b>EMFILE</b>	The <i>Command</i> parameter is <b>F_DUPFD</b> , and the maximum number of file descriptors are currently open ( <b>OPEN_MAX</b> ).
<b>EINVAL</b>	The <i>Command</i> parameter is <b>F_DUPFD</b> , and the <i>Argument</i> parameter is negative or greater than or equal to <b>OPEN_MAX</b> .
<b>EINVAL</b>	An illegal value was provided for the <i>Command</i> parameter.
<b>EINVAL</b>	An attempt was made to lock a fifo or pipe.
<b>ESRCH</b>	The value of the <i>Command</i> parameter is <b>F_SETOWN</b> , and the process ID specified as the <i>Argument</i> parameter is not in use.

- EINTR** The *Command* parameter was **F\_SETLKW** and the process received a signal while waiting to acquire the lock.
- E\_OVERFLOW** The *Command* parameter was **F\_GETLK** and the block lock could not be represented in the **flock** structure.

The **dup** and **dup2** subroutines fail if one or both of the following are true:

- EBADF** The *Old* parameter specifies an invalid open file descriptor or the *New* parameter specifies a file descriptor that is out of range.
- EMFILE** The number of file descriptors exceeds the **OPEN\_MAX** value or there is no file descriptor above the value of the *New* parameter.

If NFS is installed on the system, the **fcntl** subroutine can fail if the following is true:

- ETIMEDOUT** The connection timed out.

## Related Information

The **close** (“close Subroutine” on page 175) subroutine, **execl**, **execv**, **execle**, **execve**, **execvp**, **execvp**, or **exec** (“exec: execl, execle, execlp, execv, execve, execvp, or exec Subroutine” on page 236) subroutines, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **ioctl** or **ioctlx** (“ioctl, ioctlx, ioctl32, or ioctl32x Subroutine” on page 561) subroutine, **lockf** (“lockfx, lockf, flock, or lockf64 Subroutine” on page 740) subroutine, **open**, **openx**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutines, **read** subroutine, **write** subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fdetach Subroutine

### Purpose

Detaches STREAMS-based file from the file to which it was attached.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stropts.h>
int fdetach(const char *path);
```

### Parameters

*path* Pathname of a file previous associated with a STREAMS-based object using the **fattach** subroutine.

### Description

The **fdetach** subroutine detaches a STREAMS-based file from the file to which it was attached by a previous call to **fattach** subroutine. The *path* argument points to the pathname of the attached STREAMS file. The process must have appropriate privileges or be the owner of the file. A successful call to **fdetach** subroutine causes all pathnames that named the attached STREAMS file to again name the file to which the STREAMS file was attached. All subsequent operations on *path* will operate on the underlying file and not on the STREAMS file.

All open file descriptors established while the STREAMS file was attached to the file referenced by *path* will still refer to the STREAMS file after the **fdetach** subroutine has taken effect.

If there are no open file descriptors or other references to the STREAMS file, then a successful call to **fdetach** subroutine has the same effect as performing the last **close** subroutine on the attached file.

The **umount** command may be used to detach a file name if an I application exits before performing **fdetach** subroutine.

## Return Value

0 Successful completion.  
-1 Not successful and **errno** set to one of the following.

## Errno Value

<b>EACCES</b>	Search permission is denied on a component of the path prefix.
<b>EPERM</b>	The effective user ID is not the owner of <i>path</i> and the process does not have appropriate privileges.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>ENOENT</b>	A component of <i>path</i> parameter does not name an existing file or <i>path</i> is an empty string.
<b>EINVAL</b>	The <i>path</i> parameter names a file that is not currently attached.
<b>ENAMETOOLONG</b>	The size of <i>path</i> parameter exceeds { <b>PATH_MAX</b> }, or a component of <i>path</i> is longer than { <b>NAME_MAX</b> }.
<b>ELOOP</b>	Too many symbolic links were encountered in resolving the <i>path</i> parameter.
<b>ENOMEM</b>	Insufficient storage space is available.

## Related Information

The **fattach** (“fattach Subroutine” on page 250) subroutine, **isastream** subroutine.

---

## fdim, fdimf, fdiml, fdimd32, fdimd64, and fdimd128 Subroutines

### Purpose

Computes the positive difference between two floating-point numbers.

### Syntax

```
#include <math.h>
```

```
double fdim (x, y)
double x;
double y;
```

```
float fdimf (x, y)
float x;
float y;
```

```
long double fdiml (x, y)
long double x;
long double y;
```

```
_Decimal32 fdimd32 (x, y);
_Decimal32 x;
_Decimal32 y;
```

```
_Decimal64 fdimd64 (x, y);
```

```
_Decimal164 x;  
_Decimal164 y;  
  
_Decimal128 fdimd128 (x, y);  
_Decimal128 x;  
_Decimal128 y;
```

## Description

The **fdim**, **fdimf**, **fdiml**, **fdimd32**, **fdimd64**, and **fdimd128** subroutines determine the positive difference between their arguments. If the value of the *x* parameter is greater than that of the *y* parameter, *x* - *y* is returned. If *x* is less than or equal to *y*, +0 is returned.

An application that wants to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. On return, if the **errno** is a value of non-zero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is a value of non-zero, an error has occurred.

## Parameters

*x*                Specifies the value to be computed.  
*y*                Specifies the value to be computed.

## Return Values

Upon successful completion, the **fdim**, **fdimf**, **fdiml**, **fdimd32**, **fdimd64**, and **fdimd128** subroutines return the positive difference value.

If *x*-*y* is positive and overflows, a range error occurs and the **fdim**, **fdimf**, **fdiml**, **fdimd32**, **fdimd64**, and **fdimd128** subroutines return the value of the **HUGE\_VAL**, **HUGE\_VALF**, **HUGE\_VALL**, **HUGE\_VAL\_D32**, **HUGE\_VAL\_D64** and **HUGE\_VAL\_D128** macro respectively.

If *x*-*y* is positive and underflows, a range error might occur, and 0.0 is returned.

If *x* or *y* is NaN, a NaN is returned.

## Related Information

“**feclearexcept** Subroutine” on page 265, “**fetestexcept** Subroutine” on page 273, “**fmax**, **fmaxf**, **fmaxl**, **fmaxd32**, **fmaxd64**, and **fmaxd128** Subroutines” on page 280, and “**madd**, **msub**, **mult**, **mdiv**, **pow**, **gcd**, **invert**, **rpow**, **msqrt**, **mcmp**, **move**, **min**, **omin**, **fmin**, **m\_in**, **mout**, **omout**, **fmout**, **m\_out**, **sdiv**, or **itom** Subroutine” on page 785.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fe\_dec\_getround and fe\_dec\_setround Subroutines

### Purpose

Reads and sets the IEEE decimal floating-point rounding mode.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <fenv.h>
int fe_dec_getround ();
int fe_dec_setround (RoundMode);
int RoundMode
```

## Description

The **fe\_dec\_getround** subroutine returns the current rounding mode. The **fe\_dec\_setround** subroutine changes the rounding mode to the *RoundMode* parameter and returns the value of 0 if it successfully completes.

Decimal floating-point rounding occurs when the infinitely precise result of a decimal floating-point operation cannot be represented exactly in the destination decimal floating-point format. The IEEE Standard for Decimal Floating-Point Arithmetic allows floating-point numbers to be rounded in five different ways: **round toward zero**, **round to nearest**, **round toward +INF**, **round toward -INF**, and **round to nearest ties away from zero**. Once a rounding mode is selected, it affects all subsequent decimal floating-point operations until another rounding mode is selected.

**Tip:** The default decimal floating-point rounding mode is the **round to nearest** mode. All C main programs begin with the rounding mode that is set to **round to nearest**.

The encodings of the rounding modes are defined in the ANSI C Standard. The **fenv.h** file contains definitions for the rounding modes. The following table shows the **fenv.h** definition, the ANSI C Standard value, and a description of each rounding mode.

fenv.h definition	ANSI value	Description
FE_DEC_TONEAREST	0	Round to nearest
FE_DEC_TOWARDZERO	1	Round toward 0
FE_DEC_UPWARD	2	Round toward +INF
FE_DEC_DOWNWARD	3	Round toward -INF
FE_DEC_TONEARESTFROMZERO	4	Round to nearest ties away from zero

## Parameters

<i>RoundMode</i>	Specifies one of the following modes: FE_DEC_TOWARDZERO, FE_DEC_TONEAREST, FE_DEC_UPWARD, FE_DEC_DOWNWARD, FE_DEC_TONEARESTFROMZERO.
------------------	---

## Return Values

Upon successful completion, the **fe\_dec\_getround** subroutine returns the current rounding mode. Otherwise, it returns -1.

Upon successful completion, the **fe\_dec\_setround** subroutine returns the value of zero. Otherwise, it returns -1.

## Related Information

The “floor, floorf, floorl, floord32, floord64, floord128, nearest, trunc, itrunc, and uitrunc Subroutines” on page 277, “ceil, ceilf, ceill, ceild32, ceild64, and ceild128 Subroutines” on page 140, “fmod, fmodf, or fmodl Subroutine” on page 282, “fabsf, fabsf, fabs, fabsd32, fabsd64, and fabsd128 Subroutines” on page 249,



“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294, and the “fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine” on page 295.

---

## feclearexcept Subroutine

### Purpose

Clears floating-point exceptions.

### Syntax

```
#include <fenv.h>

int feclearexcept (excepts)
int excepts;
```

### Description

The **feclearexcept** subroutine attempts to clear the supported floating-point exceptions represented by the *excepts* parameter.

### Parameters

*excepts*                      Specifies the supported floating-point exception to be cleared.

### Return Values

If the *excepts* parameter is zero or if all the specified exceptions were successfully cleared, the **feclearexcept** subroutine returns zero. Otherwise, it returns a nonzero value.

### Related Information

“fegetexceptflag or fesetexceptflag Subroutine” on page 266, “feraiseexcept Subroutine” on page 270, and “fetestexcept Subroutine” on page 273

---

## fegetenv or fesetenv Subroutine

### Purpose

Gets and sets the current floating-point environment.

### Syntax

```
#include <fenv.h>

int fegetenv (envp)
fenv_t *envp;

int fesetenv (envp)
const fenv_t *envp;
```

### Description

The **fegetenv** subroutine stores the current floating-point environment in the object pointed to by the *envp* parameter.

The **fesetenv** subroutine attempts to establish the floating-point environment represented by the object pointed to by the *envp* parameter. The *envp* parameter points to an object set by a call to the **fegetenv** or

**feholdexcept** subroutines, or equal a floating-point environment macro. The **fesetenv** subroutine does not raise floating-point exceptions. It only installs the state of the floating-point status flags represented through its argument.

## Parameters

*envp* Points to an object set by a call to the **fegetenv** or **feholdexcept** subroutines, or equal a floating-point environment macro.

## Return Values

If the representation was successfully stored, the **fegetenv** subroutine returns zero. Otherwise, it returns a nonzero value. If the environment was successfully established, the **fesetenv** subroutine returns zero. Otherwise, it returns a nonzero value.

## Related Information

“feholdexcept Subroutine” on page 267 and “feupdateenv Subroutine” on page 274

---

## fegetexceptflag or fesetexceptflag Subroutine

### Purpose

Gets and sets floating-point status flags.

### Syntax

```
#include <fenv.h>
```

```
int fegetexceptflag (flagp, excepts)
feexcept_t *flagp;
int excepts;
```

```
int fesetexceptflag (flagp, excepts)
const feexcept_t *flagp;
int excepts;
```

### Description

The **fegetexceptflag** subroutine attempts to store an implementation-defined representation of the states of the floating-point status flags indicated by the *excepts* parameter in the object pointed to by the *flagp* parameter.

The **fesetexceptflag** subroutine attempts to set the floating-point status flags indicated by the *excepts* parameter to the states stored in the object pointed to by the *flagp* parameter. The value pointed to by the *flagp* parameter shall have been set by a previous call to the **fegetexceptflag** subroutine whose second argument represented at least those floating-point exceptions represented by the *excepts* parameter. This subroutine does not raise floating-point exceptions. It only sets the state of the flags.

### Parameters

*flagp* Points to the object that holds the implementation-defined representation of the states of the floating-point status flags.

*excepts* Points to an implementation-defined representation of the states of the floating-point status flags.

## Return Values

If the representation was successfully stored, the **fegetexceptflag** parameter returns zero. Otherwise, it returns a nonzero value. If the *excepts* parameter is zero or if all the specified exceptions were successfully set, the **fesetexceptflag** subroutine returns zero. Otherwise, it returns a nonzero value.

## Related Information

“feraiseexcept Subroutine” on page 270 and “fetestexcept Subroutine” on page 273.

---

## fegetround or fesetround Subroutine

### Purpose

Gets and sets the current rounding direction.

### Syntax

```
#include <fenv.h>

int fegetround (void)

int fesetround (round)
int round;
```

### Description

The **fegetround** subroutine gets the current rounding direction.

The **fesetround** subroutine establishes the rounding direction represented by the *round* parameter. If the *round* parameter is not equal to the value of a rounding direction macro, the rounding direction is not changed.

### Parameters

*round*            Specifies the rounding direction.

### Return Values

The **fegetround** subroutine returns the value of the rounding direction macro representing the current rounding direction or a negative value if there is no such rounding direction macro or the current rounding direction is not determinable.

The **fesetround** subroutine returns a zero value if the requested rounding direction was established.

---

## feholdexcept Subroutine

### Purpose

Saves current floating-point environment.

### Syntax

```
#include <fenv.h>

int feholdexcept (envp)
fenv_t *envp;
```

## Description

The **feholdexcept** subroutine saves the current floating-point environment in the object pointed to by *envp*, clears the floating-point status flags, and installs a non-stop (continue on floating-point exceptions) mode for all floating-point exceptions.

## Parameters

*envp*                Points to the current floating-point environment.

## Return Values

The **feholdexcept** subroutine returns zero if non-stop floating-point exception handling was successfully installed.

## Related Information

The “feupdateenv Subroutine” on page 274.

---

## fence Subroutine

### Purpose

Allows you to request and change the virtual shared disk fence map.

### Syntax

```
#include <vsd_ioctl.h>
int ioctl(FileDescriptor, Command, Argument)
int FileDescriptor, Command;
void *Argument;
```

### Description

Use this subroutine to request and change the virtual shared disk fence map. The fence map, which controls whether virtual shared disks can send or satisfy requests from virtual shared disks at remote nodes, is defined as:

```
struct vsd_FenceMap     /* This is the argument to the VSD fence ioctl. */
{
    ulong                flags;
    vsd_minorBitmap_t    minornoBitmap; /* Bitmap of minor numbers to fence
                                         (supports 10000 vsds)                */
    vsd_Fence_Bitmap_t   nodesBitmap;  /* Nodes to (un)fence these vsds from
                                         (supports node numbers 1-2048)    */
}vsd_FenceMap_t
```

The flags **VSD\_FENCE** and **VSD\_UNFENCE** are mutually exclusive — an *ioctl* can either fence a set of virtual shared disks or unfence a set of virtual shared disks, but not both. The *minornoBitmap* denotes which virtual shared disks are to be fenced/unfenced from the nodes specified in the *nodesBitmap*.

### Parameters

*FileDescriptor*    Specifies the open file descriptor for which the control operation is to be performed.

*Command*            Specifies the control function to be performed. The value of this parameter is always **GIOCFENCE**.

*Argument*            Specifies a pointer to a **vsd\_fence\_map** structure.

The *flags* field of the **vsd\_fence\_map** structure determines the type of operation that is performed. The flags could be set with one or more options using the OR operator. These options are as follows:

<b>VSD_FENCE_FORCE</b>	If this option is specified, a node can unfence itself.
<b>VSD_FENCE_GET</b>	Denotes a query request.
<b>VSD_FENCE</b>	Denotes a fence request.
<b>VSD_UNFENCE</b>	Denotes an unfence request.

## Examples

The following example fences a virtual shared disk with a minor number of 7 from node 4 and 5, and unfences a virtual shared disk with a minor number of 5 from node 1:

```
int fd;
vsd_FenceMap_t FenceMap;

/* Clear the FenceMap */
bzero(FenceMap, sizeof(vsd_FenceMap_t));

/* fence nodes 4,5 from minor 7 */
FenceMap.flags = VSD_FENCE;
MAP_SET(7, FenceMap.minornoBitmap);
MAP_SET(4, FenceMap.nodesBitmap);
MAP_SET(5, FenceMap.nodesBitmap);

/* Issue the fence request */
ioctl(fd, GIOCFENCE, &FenceMap);

/* Unfence node 1 from minor 5 */
bzero(FenceMap, sizeof(vsd_FenceMap_t));
FenceMap.flags = VSD_UNFENCE | VSD_FENCE_FORCE;
MAP_SET(5, FenceMap.minornoBitmap);
MAP_SET(1, FenceMap.nodesBitmap);

/* Issue the fence request */
ioctl(fd, GIOCFENCE, &FenceMap);
```

## Return Values

If the request succeeds, the `ioctl` returns 0. In the case of an error, a value of -1 is returned with the global variable **errno** set to identify the error.

## Error Values

The **fence** `ioctl` subroutine can return the following error codes:

<b>EACCES</b>	Indicates that an unfence was requested from a fenced node without the <b>VSD_FENCE_FORCE</b> option.
<b>EINVAL</b>	Indicates an invalid request (ambiguous flags or unidentified virtual shared disks).
<b>ENOCONNECT</b>	Indicates that either the primary or the secondary node for a virtual shared disk to be fenced is not a member of the virtual shared disk group, or the virtual shared disk in question is in the <b>stopped</b> state.
<b>ENOTREADY</b>	Indicates that the group is not active or the Recoverable virtual shared disk subsystem is not available.
<b>ENXIO</b>	Indicates that the Virtual shared disk driver is being unloaded.

---

## feof, ferror, clearerr, or fileno Macro

### Purpose

Checks the status of a stream.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
int feof ( Stream)
```

```
FILE *Stream;
```

```
int ferror (Stream)
```

```
FILE *Stream;
```

```
void clearerr (Stream)
```

```
FILE *Stream;
```

```
int fileno (Stream)
```

```
FILE *Stream;
```

### Description

The **feof** macro inquires about the end-of-file character (EOF). If EOF has previously been detected reading the input stream specified by the *Stream* parameter, a nonzero value is returned. Otherwise, a value of 0 is returned.

The **ferror** macro inquires about input or output errors. If an I/O error has previously occurred when reading from or writing to the stream specified by the *Stream* parameter, a nonzero value is returned. Otherwise, a value of 0 is returned.

The **clearerr** macro inquires about the status of a stream. The **clearerr** macro resets the error indicator and the EOF indicator to a value of 0 for the stream specified by the *Stream* parameter.

The **fileno** macro inquires about the status of a stream. The **fileno** macro returns the integer file descriptor associated with the stream pointed to by the *Stream* parameter. Otherwise a value of -1 is returned.

### Parameters

*Stream*                Specifies the input or output stream.

### Related Information

The **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine.

Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## feraiseexcept Subroutine

### Purpose

Raises the floating-point exception.

## Syntax

```
#include <fenv.h>
```

```
int feraiseexcept (excepts)  
int excepts;
```

## Description

The **feraiseexcept** subroutine attempts to raise the supported floating-point exceptions represented by the *excepts* parameter. The order in which these floating-point exceptions are raised is unspecified.

## Parameters

*excepts*        Points to the floating-point exceptions.

## Return Values

If the argument is zero or if all the specified exceptions were successfully raised, the **feraiseexcept** subroutine returns a zero. Otherwise, it returns a nonzero value.

## Related Information

“feclearexcept Subroutine” on page 265, “fegetexceptflag or fesetexceptflag Subroutine” on page 266, “fetestexcept Subroutine” on page 273.

---

## fetch\_and\_add Subroutine

### Purpose

Updates a single word variable atomically.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/atomic_op.h>
```

```
int fetch_and_add ( word_addr, value)  
atomic_p word_addr;  
int value;
```

### Description

The **fetch\_and\_add** subroutine increments one word in a single atomic operation. This operation is useful when a counter variable is shared between several threads or processes. When updating such a counter variable, it is important to make sure that the fetch, update, and store operations occur atomically (are not interruptible). For example, consider the sequence of events which could occur if the operations were interruptible:

1. A process fetches the counter value and adds one to it.
2. A second process fetches the counter value, adds one, and stores it.
3. The first process stores its value.

The result of this is that the update made by the second process is lost.

Traditionally, atomic access to a shared variable would be controlled by a mechanism such as semaphores. Compared to such mechanisms, the **fetch\_and\_add** subroutine requires very little overhead, and provided that the counter variable fits in a single machine word, this subroutine provides a highly efficient way of performing this operation.

**Note:** The word containing the counter variable must be aligned on a full word boundary.

## Parameters

<i>word_addr</i>	Specifies the address of the word variable to be incremented.
<i>value</i>	Specifies the value to be added to the word variable.

## Return Values

This subroutine returns the original value of the word.

## Related Information

The **fetch\_and\_and** (“fetch\_and\_and or fetch\_and\_or Subroutine”) subroutine, **fetch\_and\_or** (“fetch\_and\_and or fetch\_and\_or Subroutine”) subroutine, **compare\_and\_swap** (“compare\_and\_swap Subroutine” on page 176) subroutine.

---

## fetch\_and\_and or fetch\_and\_or Subroutine

### Purpose

Sets or clears bits in a single word variable atomically.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/atomic_op.h>
```

```
uint fetch_and_and ( word_addr, mask)
atomic_p word_addr;
int mask;
```

```
uint fetch_and_or ( word_addr, mask)
atomic_p word_addr;
int mask;
```

### Description

The **fetch\_and\_and** and **fetch\_and\_or** subroutines respectively clear and set bits in one word, according to a bit mask, in a single atomic operation. The **fetch\_and\_and** subroutine clears bits in the word which correspond to clear bits in the bit mask, and the **fetch\_and\_or** subroutine sets bits in the word which correspond to set bits in the bit mask.

These operations are useful when a variable containing bit flags is shared between several threads or processes. When updating such a variable, it is important that the fetch, bit clear or set, and store operations occur atomically (are not interruptible). For example, consider the sequence of events which could occur if the operations were interruptible:

1. A process fetches the flags variable and sets a bit in it.
2. A second process fetches the flags variable, sets a different bit, and stores it.



3. The first process stores its value.

The result is that the update made by the second process is lost.

Traditionally, atomic access to a shared variable would be controlled by a mechanism such as semaphores. Compared to such mechanisms, the **fetch\_and\_and** and **fetch\_and\_or** subroutines require very little overhead, and provided that the flags variable fits in a single machine word, they provide a highly efficient way of performing this operation.

**Note:** The word containing the flag bits must be aligned on a full word boundary.

## Parameters

*word\_addr* Specifies the address of the single word variable whose bits are to be cleared or set.  
*mask* Specifies the bit mask which is to be applied to the single word variable.

## Return Values

These subroutines return the original value of the word.

## Related Information

The **fetch\_and\_add** (“fetch\_and\_add Subroutine” on page 271) subroutine, **compare\_and\_swap** (“compare\_and\_swap Subroutine” on page 176) subroutine.

---

## fetestexcept Subroutine

### Purpose

Tests floating-point exception flags.

### Syntax

```
#include <fenv.h>

int fetestexcept (excepts)
int excepts;
```

### Description

The **fetestexcept** subroutine determines which of a specified subset of the floating-point exception flags are currently set. The *excepts* parameter specifies the floating-point status flags to be queried.

### Parameters

*excepts* Specifies the floating-point status flags to be queried.

### Return Values

The **fetestexcept** subroutine returns the value of the bitwise-inclusive OR of the floating-point exception macros corresponding to the currently set floating-point exceptions included in *excepts*.

### Related Information

“feclearexcept Subroutine” on page 265, “fegetexceptflag or fesetexceptflag Subroutine” on page 266, and “feraiseexcept Subroutine” on page 270

---

## feupdateenv Subroutine

### Purpose

Updates floating-point environment.

### Syntax

```
#include <fenv.h>
```

```
int feupdateenv (envp)
const fenv_t *envp;
```

### Description

The **feupdateenv** subroutine attempts to save the currently raised floating-point exceptions in its automatic storage, attempts to install the floating-point environment represented by the object pointed to by the *envp* parameter, and attempts to raise the saved floating-point exceptions. The *envp* parameter point to an object set by a call to **feholdexcept** or **fegetenv**, or equal a floating-point environment macro.

### Parameters

*envp*            Points to an object set by a call to the **feholdexcept** or the **fegetenv** subroutine, or equal a floating-point environment macro.

### Return Values

The **feupdateenv** subroutine returns a zero value if all the required actions were successfully carried out.

### Related Information

“fegetenv or fesetenv Subroutine” on page 265 and “feholdexcept Subroutine” on page 267.

---

## finfo or ffinfo Subroutine

### Purpose

Returns file information.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/finfo.h>
```

```
int finfo(Path1, cmd, buffer, length)
const char *Path1;
int cmd;
void *buffer;
int length;
```

```
int ffinfo (fd, cmd, buffer, length)
int fd;
int cmd;
void *buffer;
int length;
```

### Description

The **finfo** and **ffinfo** subroutines return specific file information for the specified file.

## Parameters

<i>Path1</i>	Path name of a file system object to query.
<i>fd</i>	File descriptor for an open file to query.
<i>cmd</i>	Specifies the type of file information to be returned.
<i>buffer</i>	User supplied buffer which contains the file information upon successful return. /usr/include/sys/finfo.h describes the buffer.
<i>length</i>	Length of the query buffer.

## Commands

### F\_PATHCONF

When the **F\_PATHCONF** command is specified, a file's implementation information is returned.

**Note:** The operating system provides another subroutine that retrieves file implementation characteristics, **pathconf** ("pathconf or fpathconf Subroutine" on page 979) command. While the **finfo** and **ffinfo** subroutines can be used to retrieve file information, it is preferred that programs use the **pathconf** interface.

### F\_DIOCAP

When the **F\_DIOCAP** command is specified, the file's direct I/O capability information is returned. The buffer supplied by the application is of type **struct diocapbuf** \*.

## Return Values

Upon successful completion, the **finfo** and **ffinfo** subroutines return a value of 0 and the user supplied buffer is correctly filled in with the file information requested. If the **finfo** or **ffinfo** subroutines were unsuccessful, a value of -1 is returned and the global **errno** variable is set to indicate the error.

## Error Codes

### EACCES

Search permission is denied for a component of the path prefix.

### EINVAL

If the length specified for the user buffer is greater than **MAX\_FINFO\_BUF**.

If the command argument is not supported. If **F\_DIOCAP** command is specified and the file object does not support Direct I/O.

### ENAMETOOLONG

The length of the Path parameter string exceeds the **PATH\_MAX** value.

### ENOENT

The named file does not exist or the Path parameter points to an empty string.

### ENOTDIR

A component of the path prefix is not a directory.

### EBADF

File descriptor provided is not valid.

## Related Information

The **pathconf** ("pathconf or fpathconf Subroutine" on page 979) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## flockfile, ftrylockfile, funlockfile Subroutine

### Purpose

Provides for explicit application-level locking of stdio (FILE\*) objects.

## Library

Standard Library (**libc.a**)

## Syntax

```
#include <stdio.h>
void flockfile (FILE * file)
int  ftrylockfile (FILE * file)
void funlockfile (FILE * file)
```

## Description

The **flockfile**, **ftrylockfile** and **funlockfile** functions provide for explicit application-level locking of **stdio** (**FILE\***) objects. These functions can be used by a thread to delineate a sequence of I/O statements that are to be executed as a unit.

The **flockfile** function is used by a thread to acquire ownership of a (**FILE\***) object.

The **ftrylockfile** function is used by a thread to acquire ownership of a (**FILE\***) object if the object is available; **ftrylockfile** is a non-blocking version of **flockfile**.

The **funlockfile** function is used to relinquish the ownership granted to the thread. The behavior is undefined if a thread other than the current owner calls the **funlockfile** function.

Logically, there is a lock count associated with each (**FILE\***) object. This count is implicitly initialised to zero when the (**FILE\***) object is created. The (**FILE\***) object is unlocked when the count is zero. When the count is positive, a single thread owns the (**FILE\***) object. When the **flockfile** function is called, if the count is zero or if the count is positive and the caller owns the (**FILE\***) object, the count is incremented. Otherwise, the calling thread is suspended, waiting for the count to return to zero. Each call to **funlockfile** decrements the count. This allows matching calls to **flockfile** (or successful calls to **ftrylockfile**) and **funlockfile** to be nested.

All functions that reference (**FILE\***) objects behave as if they use **flockfile** and **funlockfile** internally to obtain ownership of these (**FILE\***) objects.

## Return Values

None for **flockfile** and **funlockfile**. The function **ftrylock** returns zero for success and non-zero to indicate that the lock cannot be acquired.

## Implementation Specifics

These subroutines are part of Base Operating System (BOS) subroutines.

Realtime applications may encounter priority inversion when using FILE locks. The problem occurs when a high priority thread locks a file that is about to be unlocked by a low priority thread, but the low priority thread is preempted by a medium priority thread. This scenario leads to priority inversion; a high priority thread is blocked by lower priority threads for an unlimited period of time. During system design, realtime programmers must take into account the possibility of this kind of priority inversion. They can deal with it in a number of 7434 ways, such as by having critical sections that are guarded by file locks execute at a high priority, so that a thread cannot be preempted while executing in its critical section.

## Related Information

The **getc\_unlocked**, **getchar\_unlocked**, **putc\_unlocked**, **putchar\_unlocked** (“**getc\_unlocked**, **getchar\_unlocked**, **putc\_unlocked**, **putchar\_unlocked** Subroutines” on page 349) subroutine.

---

## floor, floorf, floorl, floord32, floord64, floord128, nearest, trunc, itrunc, and uitrunc Subroutines

### Purpose

The **floor** subroutine, **floorf** subroutine, **floorl** subroutine, **nearest** subroutine, **trunc** subroutine, **floord32** subroutine, **floord64** subroutine, and **floord128** subroutine, round floating-point numbers to floating-point integer values.

The **itrunc** subroutine and **uitrunc** subroutine round floating-point numbers to signed and unsigned integers, respectively.

### Libraries

IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)  
Standard C Library (**libc.a**) (separate syntax follows)

### Syntax

```
#include <math.h>
```

```
double floor ( x )  
double x;
```

```
float floorf ( x )  
float x;
```

```
long double floorl ( x )  
long double x;
```

```
_Decimal32 floord32(x)  
_Decimal32 x;
```

```
_Decimal64 floord64(x)  
_Decimal64 x;
```

```
_Decimal128 floord128(x)  
_Decimal128 x;
```

```
double nearest ( x )  
double x;
```

```
double trunc ( x )  
double x;
```

Standard C Library (**libc.a**)

```
#include <stdlib.h>  
#include <limits.h>
```

```
int itrunc ( x )  
double x;
```

```
unsigned int uitrunc ( x )  
double x;
```

## Description

The **floor**, **floorf**, **floorl**, **floord32**, **floord64**, and **floord128** subroutines return the largest floating-point integer value that is not greater than the *x* parameter.

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

The **nearest** subroutine returns the nearest floating-point integer value to the *x* parameter. If *x* lies exactly halfway between the two nearest floating-point integer values, an even floating-point integer is returned.

The **trunc** subroutine returns the nearest floating-point integer value to the *x* parameter in the direction of 0. This is equivalent to truncating off the fraction bits of the *x* parameter.

**Note:** The default floating-point rounding mode is *round to nearest*. All C main programs begin with the rounding mode set to *round to nearest*.

The **itrunc** subroutine returns the nearest signed integer to the *x* parameter in the direction of 0. This is equivalent to truncating the fraction bits from the *x* parameter and then converting *x* to a signed integer.

The **uitrunc** subroutine returns the nearest unsigned integer to the *x* parameter in the direction of 0. This action is equivalent to truncating off the fraction bits of the *x* parameter and then converting *x* to an unsigned integer.

**Note:** Compile any routine that uses subroutines from the **libm.a** library with the **-lm** flag. To compile the `floor.c` file, for example, enter:

```
cc floor.c -lm
```

The **itrunc**, **uitrunc**, **trunc**, and **nearest** subroutines are not part of the ANSI C Library.

## Parameters

*x* Specifies a double-precision floating-point value. For the **floorl** subroutine, specifies a long double-precision floating-point value.

*y* Specifies a double-precision floating-point value. For the **floorl** subroutine, specifies some long double-precision floating-point value.

## Return Values

Upon successful completion, the **floor**, **floorf**, **floorl**, **floord32**, **floord64**, and **floord128** subroutines return the largest integral value that is not greater than *x*, expressed as a **double**, **float**, **long double**, **\_Decimal32**, **\_Decimal64**, or **\_Decimal128**, as appropriate for the return type of the function.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* is returned.

If the correct value would cause overflow, a range error occurs and the **floor**, **floorf**, **floorl**, **floord32**, **floord64**, and **floord128** subroutines return the value of the macro **-HUGE\_VAL**, **-HUGE\_VALF**, **-HUGE\_VALL**, **-HUGE\_VAL\_D32**, **-HUGE\_VAL\_D64**, and **-HUGE\_VAL\_D128**, respectively.

## Error Codes

The **itrunc** and **utrunc** subroutines return the **INT\_MAX** value if  $x$  is greater than or equal to the **INT\_MAX** value and the **INT\_MIN** value if  $x$  is equal to or less than the **INT\_MIN** value. The **itrunc** subroutine returns the **INT\_MIN** value if  $x$  is a Quiet NaN(not-a-number) or Silent NaN. The **utrunc** subroutine returns 0 if  $x$  is a Quiet NaN or Silent NaN. (The **INT\_MAX** and **INT\_MIN** values are defined in the **limits.h** file.) The **utrunc** subroutine **INT\_MAX** if  $x$  is greater than **INT\_MAX** and 0 if  $x$  is less than or equal 0.0

## Files

**float.h**            Contains the ANSI C **FLT\_ROUNDS** macro.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

The **fp\_read\_rnd** or **fp\_swap\_rnd** (“fp\_read\_rnd or fp\_swap\_rnd Subroutine” on page 302) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fma, fmaf, or fmal Subroutine

### Purpose

Floating-point multiply-add.

### Syntax

```
#include <math.h>
```

```
double fma (x, y, z)
double x;
double y;
double z;
```

```
float fmaf (x, y, z)
float x;
float y;
float z;
```

```
long double fmal (x, y, z)
long double x;
long double y;
long double z;
```

### Description

The **fma**, **fmaf**, and **fmal** subroutines compute  $(x * y) + z$ , rounded as one ternary operation. They compute the value (as if) to infinite precision and round once to the result format, according to the rounding mode characterized by the value of **FLT\_ROUNDS**.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be multiplied by the *y* parameter.  
*y* Specifies the value to be multiplied by the *x* parameter.  
*z* Specifies the value to be added to the product of the *x* and *y* parameters.

## Return Values

Upon successful completion, the **fma**, **fmaf**, and **fmal** subroutines return  $(x * y) + z$ , rounded as one ternary operation.

If *x* or *y* are NaN, a NaN is returned.

If *x* multiplied by *y* is an exact infinity and *z* is also an infinity but with the opposite sign, a domain error occurs, and a NaN is returned.

If one of the *x* and *y* parameters is infinite, the other is zero, and the *z* parameter is not a NaN, a domain error occurs, and a NaN is returned.

If one of the *x* and *y* parameters is infinite, the other is zero, and *z* is a NaN, a NaN is returned and a domain error may occur.

If  $x*y$  is not  $0*Inf$  nor  $Inf*0$  and *z* is a NaN, a NaN is returned.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fmax, fmaxf, fmaxl, fmaxd32, fmaxd64, and fmaxd128 Subroutines

### Purpose

Determines the maximum numeric value of two floating-point numbers.

### Syntax

```
#include <math.h>
```

```
double fmax (x, y)
double x;
double y;
```

```
float fmaxf (x, y)
float x;
float y;
```

```
long double fmaxl (x, y)
long double x;
long double y;
```

```
_Decimal32 fmaxd32 (x, y);
_Decimal32 x;
```



```

_Decimal132 y;

_Decimal164 fmaxd64 (x, y);
_Decimal164 x;
_Decimal164 y;

_Decimal128 fmaxd128 (x, y);
_Decimal128 x;
_Decimal128 y;

```

## Description

The **fmax**, **fmaxf**, **fmaxl**, **fmaxd32**, **fmaxd64**, and **fmaxd128** subroutines determine the maximum numeric value of their arguments. NaN arguments are treated as missing data. If one argument is a NaN and the other numeric, the **fmax**, **fmaxf**, **fmaxl**, **fmaxd32**, **fmaxd64**, and **fmaxd128** subroutines choose the numeric value.

## Parameters

*x*                Specifies the value to be computed.  
*y*                Specifies the value to be computed.

## Return Values

Upon successful completion, the **fmaxl**, **fmaxf**, **fmaxl**, **fmaxd32**, **fmaxd64**, and **fmaxd128** subroutines return the maximum numeric value of their arguments.

If one argument is a NaN, the other argument is returned.

If *x* and *y* are NaN, a NaN is returned.

## Related Information

“**fdim**, **fdimf**, **fdiml**, **fdimd32**, **fdimd64**, and **fdimd128** Subroutines” on page 262 and “**madd**, **msub**, **mult**, **mdiv**, **pow**, **gcd**, **invert**, **rpow**, **msqrt**, **mcmp**, **move**, **min**, **omin**, **fmin**, **m\_in**, **mout**, **omout**, **fmout**, **m\_out**, **sdiv**, or **itom** Subroutine” on page 785

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fminf, fminl, fmind32, fmind64, and fmind128 Subroutines

### Purpose

Determines the minimum numeric value of two floating-point numbers.

### Syntax

```
#include <math.h>
```

```
float fminf (x, y)
float x;
float y;
```

```
long double fminl (x, y)
long double x;
long double y;
```

```
_Decimal132 fmind32 (x, y)
_Decimal132 x;
_Decimal132 y;
```

```
_Decimal64 fmind64 (x, y)
_Decimal64 x;
_Decimal64 y;
```

```
_Decimal128 fmind128 (x, y)
_Decimal128 x;
_Decimal128 y;
```

## Description

The **fminf**, **fminl**, **fmind32**, **fmind64**, and **fmind128** subroutines determine the minimum numeric value of their arguments. NaN arguments are treated as missing data. If one argument is a NaN and the other numeric, the **fminf**, **fminl**, **fmind32**, **fmind64**, and **fmind128** subroutines choose the numeric value.

## Parameters

*x* Specifies the value to be computed.  
*y* Specifies the value to be computed.

## Return Values

Upon successful completion, the **fminf**, **fminl**, **fmind32**, **fmind64**, and **fmind128** subroutines return the minimum numeric value of their arguments.

If one argument is a NaN, the other argument is returned.

If *x* and *y* are NaN, a NaN is returned.

## Related Information

“fdim, fdimf, fdiml, fdimd32, fdimd64, and fdimd128 Subroutines” on page 262 and “fmax, fmaxf, fmaxl, fmaxd32, fmaxd64, and fmaxd128 Subroutines” on page 280.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fmod, fmodf, or fmodl Subroutine

### Purpose

Computes the floating-point remainder value.

### Syntax

```
#include <math.h>
```

```
float fmodf (x, y)
float x;
float y;
```

```
long double fmodl (x, y)
long double x, y;
```

```
double fmod (x, y)
double x, y;
```

### Description

The **fmodf**, **fmodl**, and **fmod** subroutines return the floating-point remainder of the division of *x* by *y*.

An application that wants to check for error situations must set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. On return, if **errno** is the value of non-zero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is the value of non-zero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.  
*y* Specifies the value to be computed.

## Return Values

The **fmodf**, **fmodl**, and **fmod** subroutines return the value  $x - i * y$ . For the integer *i* such that, if *y* is nonzero, the result has the same sign as *x* and the magnitude is less than the magnitude of *y*.

If the correct value will cause underflow, and is not representable, a range error might occur, and 0.0 is returned.

If *x* or *y* is NaN, a NaN is returned.

If *y* is zero, a domain error occurs, and a NaN is returned.

If *x* is infinite, a domain error occurs, and a NaN is returned.

If *x* is  $\pm 0$  and *y* is not zero,  $\pm 0$  is returned.

If *x* is not infinite and *y* is  $\pm \text{Inf}$ , *x* is returned.

If the correct value will cause underflow, and is representable, a range error might occur and the correct value is returned.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fmtmsg Subroutine

### Purpose

Display a message in the specified format on standard error, the console, or both.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fmtmsg.h>
```

```
int fmtmsg (long Classification,  
const char *Label,  
int Severity,  
const char *Text;  
const char *Action,  
const char *Tag)
```

## Description

The **fmtmsg** subroutine can be used to display messages in a specified format instead of the traditional **printf** subroutine interface.

Base on a message's classification component, the **fmtmsg** subroutine either writes a formatted message to standard error, the console, or both.

A formatted message consists of up to five parameters. The *Classification* parameter is not part of a message displayed to the user, but defines the source of the message and directs the display of the formatted message.

## Parameters

<i>Classification</i>	<p>Contains identifiers from the following groups of major classifications and subclassifications. Any one identifier from a subclass may be used in combination with a single identifier from a different subclass. Two or more identifiers from the same subclass should not be used together, with the exception of identifiers from the display subclass. (Both display subclass identifiers may be used so that messages can be displayed to both standard error and system console).</p> <p><b>major classifications</b> Identifies the source of the condition. Identifiers are: <b>MM_HARD</b> (hardware), <b>MM_SOFT</b> (software), and <b>MM_FIRM</b> (firmware).</p> <p><b>message source subclassifications</b> Identifies the type of software in which the problem is detected. Identifiers are: <b>MM_APPL</b> (application), <b>MM_UTIL</b> (utility), and <b>MM_OPSYS</b> (operating system).</p> <p><b>display subclassification</b> Indicates where the message is to be displayed. Identifiers are: <b>MM_PRINT</b> to display the message on the standard error stream, <b>MM_CONSOLE</b> to display the message on the system console. One or both identifiers may be used.</p> <p><b>status subclassifications</b> Indicates whether the application will recover from the condition. Identifiers are: <b>MM_RECOVER</b> (recoverable) and <b>MM_RECOV</b> (non-recoverable).</p> <p>An additional identifier, <b>MM_NULLMC</b>, identifies that no classification component is supplied for the message.</p>
<i>Label</i>	Identifies the source to the message. The format is two fields separated by a colon. The first field is up to 10 bytes, the second field is up to 14 bytes.
<i>Severity</i>	
<i>Text</i>	Describes the error condition that produced the message. The character string is not limited to a specific size. If the character string is null then a message will be issued stating that no text has been provided.
<i>Action</i>	Describes the first step to be taken in the error-recovery process. The <b>fmtmsg</b> subroutine precedes the action string with the prefix: T0 FIX:. The <i>Action</i> string is not limited to a specific size.
<i>Tag</i>	An identifier which references online documentation for the message. Suggested usage is that <i>tag</i> includes the <i>Label</i> and a unique identifying number. A sample <i>tag</i> is UX:cat:146.

## Environment Variables

The **MSGVERB** (message verbosity) environment variable controls the behavior of the **fmtmsg** subroutine.

**MSGVERB** tells the **fmtmsg** subroutine which message components it is to select when writing messages to standard error. The value of **MSGVERB** is a colon-separated list of optional keywords. **MSGVERB** can be set as follows:

```
MSGVERB=[keyword[:keyword[:...]]]
export MSGVERB
```

Valid keywords are: *Label*, *Severity*, *Text*, *Action*, and *Tag*. If **MSGVERB** contains a keyword for a component and the component's value is not the component's null value, **fmtmsg** subroutine includes that component in the message when writing the message to standard error. If **MSGVERB** does not include a keyword for a message component, that component is not included in the display of the message. The keywords may appear in any order. If **MSGVERB** is not defined, if its value is the null string, if its value is not of the correct format, or if it contains keywords other than the valid ones listed previously, the **fmtmsg** subroutine selects all components.

**MSGVERB** affects only which components are selected for display to standard error. All message components are included in console messages.

## Application Usage

One or more message components may be systematically omitted from messages generated by an application by using the null value of the parameter for that component. The table below indicates the null values and identifiers for **fmtmsg** subroutine parameters. The parameters are of type **char\*** unless otherwise indicated.

Parameter	Null-Value	Identifier
<i>label</i>	(char*)0	MM_NULLLBL
<i>severity</i> (type <b>int</b> )	0	MM_NULLSEV
<i>class</i> (type <b>long</b> )	0L	MM_NULLMC
<i>text</i>	(char*)0	MM_NULLTXT
<i>action</i>	(char*)0	MM_NULLACT
<i>tag</i>	(char*)0	MM_NULLTAG

Another means of systematically omitting a component is by omitting the component keywords when defining the **MSGVERB** environment variable.

## Return Values

The exit codes for the **fmtmsg** subroutine are the following:

<b>MM_OK</b>	The function succeeded.
<b>MM_NOTOK</b>	The function failed completely.
<b>MM_MOMSG</b>	The function was unable to generate a message on standard error.
<b>MM_NOCON</b>	The function was unable to generate a console message.

## Examples

- The following example of the **fmtmsg** subroutine:

```
fmtmsg(MM_PRINT, "UX:cat", MM_ERROR, "illegal option",
"refer tp cat in user's reference manual", "UX:cat:001")
```

produces a complete message in the specified message format:

```
UX:cat ERROR: illegal option
TO FIX: refer to cat in user's reference manual UX:cat:001
```

- When the environment variable **MSGVERB** is set as follows:

```
MSGVERB=severity:text:action
```

and the Example 1 is used, the **fmtmsg** subroutine produces:

ERROR: illegal option  
TO FIX: refer to cat in user's reference manual UX:cat:001

## Related Information

The **printf** (“printf, fprintf, sprintf, snprintf, wprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine.

---

## fnmatch Subroutine

### Purpose

Matches file name patterns.

### Library

Standard C Library (**libc. a**)

### Syntax

```
#include <fnmatch.h>
```

```
int fnmatch ( Pattern, String, Flags);  
int Flags;  
const char *Pattern, *String;
```

### Description

The **fnmatch** subroutine checks the string specified by the *String* parameter to see if it matches the pattern specified by the *Pattern* parameter.

The **fnmatch** subroutine can be used by an application or command that needs to read a dictionary and apply a pattern against each entry; the **find** command is an example of this. It can also be used by the **pax** command to process its *Pattern* variables, or by applications that need to match strings in a similar manner.

### Parameters

<i>Pattern</i>	Contains the pattern to which the <i>String</i> parameter is to be compared. The <i>Pattern</i> parameter can include the following special characters: <ul style="list-style-type: none"><li><b>* (asterisk)</b> Matches zero, one, or more characters.</li><li><b>? (question mark)</b> Matches any single character, but will not match 0 (zero) characters.</li><li><b>[ ] (brackets)</b> Matches any one of the characters enclosed within the brackets. If a pair of characters separated by a dash are contained within the brackets, the pattern matches any character that lexically falls between the two characters in the current locale.</li></ul>
<i>String</i>	Contains the string to be compared against the <i>Pattern</i> parameter.

*Flags* Contains a bit flag specifying the configurable attributes of the comparison to be performed by the **fnmatch** subroutine.

The *Flags* parameter modifies the interpretation of the *Pattern* and *String* parameters. It is the bitwise inclusive OR of zero or more of the following flags (defined in the **fnmatch.h** file):

**FNM\_PATHNAME**

Indicates the / (slash) in the *String* parameter matches a / in the *Pattern* parameter.

**FNM\_PERIOD**

Indicates a leading period in the *String* parameter matches a period in the *Pattern* parameter.

**FNM\_NOESCAPE**

Enables quoting of special characters using the \ (backslash).

**FNM\_IGNORECASE**

Ignores uppercase and lowercase when matching alphabetic characters (available only in AIX 5.1 or later).

If the **FNM\_PATHNAME** flag is set in the *Flags* parameter, a / (slash) in the *String* parameter is explicitly matched by a / in the *Pattern* parameter. It is not matched by either the \* (asterisk) or ? (question-mark) special characters, nor by a bracket expression. If the **FNM\_PATHNAME** flag is not set, the / is treated as an ordinary character.

If the **FNM\_PERIOD** flag is set in the *Flags* parameter, then a leading period in the *String* parameter only matches a period in the *Pattern* parameter; it is not matched by either the asterisk or question-mark special characters, nor by a bracket expression. The setting of the **FNM\_PATHNAME** flag determines a period to be leading, according to the following rules:

- If the **FNM\_PATHNAME** flag is set, a . (period) is leading only if it is the first character in the *String* parameter or if it immediately follows a /.
- If the **FNM\_PATHNAME** flag is not set, a . (period) is leading only if it is the first character of the *String* parameter. If **FNM\_PERIOD** is not set, no special restrictions are placed on matching a period.

If the **FNM\_NOESCAPE** flag is not set in the *Flags* parameter, a \ (backslash) character in the *Pattern* parameter, followed by any other character, will match that second character in the *String* parameter. For example, \\ will match a backslash in the *String* parameter. If the **FNM\_NOESCAPE** flag is set, a \ (backslash) will be treated as an ordinary character.

## Return Values

If the value in the *String* parameter matches the pattern specified by the *Pattern* parameter, the **fnmatch** subroutine returns 0. If there is no match, the **fnmatch** subroutine returns the **FNM\_NOMATCH** constant, which is defined in the **fnmatch.h** file. If an error occurs, the **fnmatch** subroutine returns a nonzero value.

## Files

**/usr/include/fnmatch.h** Contains system-defined flags and constants.

## Related Information

The **glob** (“glob Subroutine” on page 482) subroutine, **globfree** (“globfree Subroutine” on page 485) subroutine, **regcomp** subroutine, **regfree** subroutine, **regerror** subroutine, **regex** subroutine.

The **find** command, **pax** command.

## fopen, fopen64, freopen, freopen64 or fdopen Subroutine

### Purpose

Opens a stream.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
FILE *fopen ( Path, Type)  
const char *Path, *Type;
```

```
FILE *fopen64 ( Path, Type)  
char *Path, *Type;
```

```
FILE *freopen (Path, Type, Stream)  
const char *Path, *Type;  
FILE *Stream;
```

```
FILE *freopen64 (Path, Type, Stream)  
char *Path, *Type;  
FILE *Stream;
```

```
FILE *fdopen ( FileDescriptor, Type)  
int FileDescriptor;  
const char *Type;
```

### Description

The **fopen** and **fopen64** subroutines open the file named by the *Path* parameter and associate a stream with it and return a pointer to the **FILE** structure of this stream.

When you open a file for update, you can perform both input and output operations on the resulting stream. However, an output operation cannot be directly followed by an input operation without an intervening **fflush** subroutine call or a file positioning operation (**fseek**, **fseeko**, **fseeko64**, **fsetpos**, **fsetpos64** or **rewind** subroutine). Also, an input operation cannot be directly followed by an output operation without an intervening flush or file positioning operation, unless the input operation encounters the end of the file.

When you open a file for appending (that is, when the *Type* parameter is set to **a**), it is impossible to overwrite information already in the file.

If two separate processes open the same file for append, each process can write freely to the file without destroying the output being written by the other. The output from the two processes is intermixed in the order in which it is written to the file.

**Note:** If the data is buffered, it is not actually written until it is flushed.



The **freopen** and **freopen64** subroutines first attempt to flush the stream and close any file descriptor associated with the *Stream* parameter. Failure to flush the stream or close the file descriptor is ignored.

The **freopen** and **freopen64** subroutines substitute the named file in place of the open stream. The original stream is closed regardless of whether the subsequent open succeeds. The **freopen** and **freopen64** subroutines returns a pointer to the **FILE** structure associated with the *Stream* parameter. The **freopen** and **freopen64** subroutines is typically used to attach the pre-opened streams associated with standard input (**stdin**), standard output (**stdout**), and standard error (**stderr**) streams to other files.

The **fdopen** subroutine associates a stream with a file descriptor obtained from an **openx** subroutine, **dup** subroutine, **creat** subroutine, or **pipe** subroutine. These subroutines open files but do not return pointers to **FILE** structures. Many of the standard I/O package subroutines require pointers to **FILE** structures.

The *Type* parameter for the **fdopen** subroutine specifies the mode of the stream, such as **r** to open a file for reading, or **a** to open a file for appending (writing at the end of the file). The mode value of the *Type* parameter specified with the **fdopen** subroutine must agree with the mode of the file specified when the file was originally opened or created.

**Note:** Using the **fdopen** subroutine with a file descriptor obtained from a call to the **shm\_open** subroutine must be avoided and might result in an error on the next **fread**, **fwrite** or **fflush** call.

The largest value that can be represented correctly in an object of type `off_t` will be established as the offset maximum in the open file description.

## Parameters

*Path*

Points to a character string that contains the name of the file to be opened.

*Type*

Points to a character string that has one of the following values:

**r** Opens a text file for reading.

**w** Creates a new text file for writing, or opens and truncates a file to 0 length.

**a** Appends (opens a text file for writing at the end of the file, or creates a file for writing).

**rb** Opens a binary file for reading.

**wb** Creates a binary file for writing, or opens and truncates a file to 0.

**ab** Appends (opens a binary file for writing at the end of the file, or creates a file for writing).

**r+** Opens a file for update (reading and writing).

**w+** Truncates or creates a file for update.

**a+** Appends (opens a text file for writing at end of file, or creates a file for writing).

**r+b , rb+**

Opens a binary file for update (reading and writing).

**w+b , wb+**

Creates a binary file for update, or opens and truncates a file to 0 length.

**a+b , ab+**

Appends (opens a binary file for update, writing at the end of the file, or creates a file for writing).

**Note:** The operating system does not distinguish between text and binary files. The **b** value in the *Type* parameter value is ignored.

*Stream*

Specifies the input stream.

*FileDescriptor*

Specifies a valid open file descriptor.

## Return Values

If the **fdopen**, **fopen**, **fopen64**, **freopen** or **freopen64** subroutine is unsuccessful, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **fopen**, **fopen64**, **freopen** and **freopen64** subroutines are unsuccessful if the following is true:

<b>EACCES</b>	Search permission is denied on a component of the path prefix, the file exists and the permissions specified by the mode are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.
<b>ELOOP</b>	Too many symbolic links were encountered in resolving path.
<b>EINTR</b>	A signal was received during the process.
<b>EISDIR</b>	The named file is a directory and the process does not have write access to it.
<b>ENAMETOOLONG</b>	The length of the filename exceeds <b>PATH_MAX</b> or a pathname component is longer than <b>NAME_MAX</b> .
<b>EMFILE</b>	The maximum number of files allowed are currently open.
<b>ENOENT</b>	The named file does not exist or the <i>File Descriptor</i> parameter points to an empty string.
<b>ENOSPC</b>	The file is not yet created and the directory or file system to contain the new file cannot be expanded.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>ENXIO</b>	The named file is a character- or block-special file, and the device associated with this special file does not exist.
<b>EOVERFLOW</b>	The named file is a regular file and the size of the file cannot be represented correctly in an object of type <i>off_t</i> .
<b>EROFS</b>	The named file resides on a read-only file system and does not have write access.
<b>ETXTBSY</b>	The file is a pure-procedure (shared-text) file that is being executed and the process does not have write access.

The **fdopen**, **fopen**, **fopen64**, **freopen** and **freopen64** subroutines are unsuccessful if the following is true:

<b>EINVAL</b>	The value of the <i>Type</i> argument is not valid.
<b>EINVAL</b>	The value of the <i>mode</i> argument is not valid.
<b>EMFILE</b>	<b>FOPEN_MAX</b> streams are currently open in the calling process.
<b>EMFILE</b>	<b>STREAM_MAX</b> streams are currently open in the calling process.
<b>ENAMETOOLONG</b>	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <b>PATH_MAX</b> .
<b>ENOMEM</b>	Insufficient storage space is available.

The **freopen** and **fopen** subroutines are unsuccessful if the following is true:

<b>EOVERFLOW</b>	The named file is a size larger than 2 Gigabytes.
------------------	---

The **fdopen** subroutine is unsuccessful if the following is true:

<b>EBADF</b>	The value of the <i>File Descriptor</i> parameter is not valid.
--------------	---

## POSIX

<b>w</b>	Truncates to 0 length or creates text file for writing.
<b>w+</b>	Truncates to 0 length or creates text file for update.
<b>a</b>	Opens or creates text file for writing at end of file.
<b>a+</b>	Opens or creates text file for update, writing at end of file.

## SAA

At least eight streams, including three standard text streams, can open simultaneously. Both binary and text modes are supported.

## Related Information

The **fclose** or **fflush** (“fclose or fflush Subroutine” on page 253) subroutine, **fseek**, **fseeko**, **fseeko64**, **rewind**, **ftell**, **ftello**, **ftello64**, **fgetpos**, **fgetpos64** or **fsetpos** (“fseek, fseeko, fseeko64, rewind, ftell, ftello, ftello64, fgetpos, fgetpos64, fsetpos, or fsetpos64 Subroutine” on page 318) subroutine, **open**, **open64**, **openx**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **setbuf**, **setvbuf**, **setbuffer**, or **setlinebuf** subroutine.

The Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fork, f\_fork, or vfork Subroutine

### Purpose

Creates a new process.

### Libraries

**fork**, **f\_fork**, and **vfork**: Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
pid_t fork(void)
pid_t f_fork(void)
int vfork(void)
```

### Description

The **fork** subroutine creates a new process. The new process (child process) is an almost exact copy of the calling process (parent process). The child process inherits the following attributes from the parent process:

- Environment
- Close-on-exec flags (described in the **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutine)
- Signal handling settings (such as the **SIG\_DFL** value, the **SIG\_IGN** value, and the *Function Address* parameter)
- Set user ID mode bit
- Set group ID mode bit
- Profiling on and off status
- Nice value
- All attached shared libraries
- Process group ID
- **tty** group ID (described in the **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243), **atexit**, or **\_exit** subroutine, **signal** subroutine, and **raise** subroutine)
- Current directory
- Root directory

- File-mode creation mask (described in the **umask** subroutine)
- File size limit (described in the **ulimit** subroutine)
- Attached shared memory segments (described in the **shmat** subroutine)
- Attached mapped file segments (described in the **shmat** subroutine)
- Debugger process ID and multiprocess flag if the parent process has multiprocess debugging enabled (described in the **ptrace** (“ptrace, ptracex, ptrace64 Subroutine” on page 1307) subroutine).

The child process differs from the parent process in the following ways:

- The child process has only one user thread; it is the one that called the **fork** subroutine.
- The child process has a unique process ID.
- The child process ID does not match any active process group ID.
- The child process has a different parent process ID.
- The child process has its own copy of the file descriptors for the parent process. However, each file descriptor of the child process shares a common file pointer with the corresponding file descriptor of the parent process.
- All **semadj** values are cleared. For information about **semadj** values, see the **semop** subroutine.
- Process locks, text locks, and data locks are not inherited by the child process. For information about locks, see the **plock** (“plock Subroutine” on page 1026) subroutine.
- If multiprocess debugging is turned on, the **trace** flags are inherited from the parent; otherwise, the **trace** flags are reset. For information about request 0, see the **ptrace** (“ptrace, ptracex, ptrace64 Subroutine” on page 1307) subroutine.
- The child process **utime**, **stime**, **cutime**, and **cstime** subroutines are set to 0. (For more information, see the **getrusage** (“getrusage, getrusage64, times, or vtimes Subroutine” on page 428), **times**, and **vtimes** subroutines.)
- Any pending alarms are cleared in the child process. (For more information, see the **incinterval** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386), **setitimer** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386), and **alarm** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) subroutines.)
- The set of signals pending for the child process is initialized to the empty set.
- The child process can have its own copy of the message catalogue for the parent process.
- The set of signals pending for the child process is initialized as an empty set.

**Attention:** If you are using the **fork** or **vfork** subroutines with an Enhanced X-Windows, X Toolkit, or Motif application, open a separate display connection (socket) for the forked process. If the child process uses the same display connection as the parent, the X Server will not be able to interpret the resulting data.

The **f\_fork** subroutine is similar to **fork**, except for:

- It is required that the child process calls one of the **exec** functions immediately after it is created. Since the **fork** handlers are never called, the application data, mutexes and the locks are all undefined in the child process.

The **vfork** subroutine is supported as a compatibility interface for older Berkeley Software Distribution (BSD) system programs and can be used by compiling with the Berkeley Compatibility Library (**libbsd.a**).

In the Version 4 of the operating system, the parent process does not have to wait until the child either exits or executes, as it does in BSD systems. The child process is given a new address space, as in the **fork** subroutine. The child process does not share any parent address space.

**Attention:** When using the **fork** or **vfork** subroutines with an Enhanced X-Windows, X Toolkit, or Motif application, a separate display connection (socket) should be opened for the forked process. The child process should never use the same display connection as the parent. Display connections are embodied with sockets, and sockets are inherited by the child process. Any attempt to have multiple processes writing to the same display connection results in the random interleaving of X protocol packets at the word level. The resulting data written to the socket will not be valid or undefined X protocol packets, and the X Server will not be able to interpret it.

**Attention:** Although the **fork** and **vfork** subroutine may be used with Graphics Library applications, the child process must not make any additional Graphics Library subroutine calls. The child application inherits some, but not all of the graphics hardware resources of the parent. Drawing by the child process may hang the graphics adapter, the Enhanced X Server, or may cause unpredictable results and place the system into an unpredictable state.

For additional information, see the `/usr/lpp/GL/README` file.

## Return Values

Upon successful completion, the **fork** subroutine returns a value of 0 to the child process and returns the process ID of the child process to the parent process. Otherwise, a value of -1 is returned to the parent process, no child process is created, and the **errno** global variable is set to indicate the error.

## Error Codes

The **fork** subroutine is unsuccessful if one or more of the following are true:

<b>EAGAIN</b>	Exceeds the limit on the total number of processes running either systemwide or by a single user, or the system does not have the resources necessary to create another process.
<b>ENOMEM</b>	Not enough space exists for this process.
<b>EPROCLIM</b>	If WLM is running, the limit on the number of processes or threads in the class may have been met.

## Related Information

The “getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386, “bindprocessor Subroutine” on page 120, “exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236, “exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243, “getrusage, getrusage64, times, or vtimes Subroutine” on page 428, “getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386, “getpriority, setpriority, or nice Subroutine” on page 412, “plock Subroutine” on page 1026, “pthread\_atfork Subroutine” on page 1208, “ptrace, ptracex, ptrace64 Subroutine” on page 1307, **raise** subroutine, **semop** subroutine, “getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386, **shmat** subroutine, **setpriority** or **getpriority** (“getpriority, setpriority, or nice Subroutine” on page 412) subroutine, **sigaction**, **sigvec**, or **signal** subroutine, **ulimit** subroutine, **umask** subroutine, **wait**, **waitpid**, or **wait3** subroutine.

The “posix\_spawn or posix\_spawnnp Subroutine” on page 1149.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Process Duplication and Termination in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*LK provides more information about forking a multi-threaded process.

---

## fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine

### Purpose

These subroutines allow operations on the floating-point trap control.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fptrap.h>

int fp_any_enable()
int fp_is_enabled( Mask)
fptrap_t Mask;

void fp_enable_all()
void fp_enable(Mask)
fptrap_t Mask;

void fp_disable_all()
void fp_disable(Mask)
fptrap_t Mask;
```

### Description

Floating point traps must be enabled before traps can be generated. These subroutines aid in manipulating floating-point traps and identifying the trap state and type.

In order to take traps on floating point exceptions, the **fp\_trap** subroutine must first be called to put the process in serialized state, and the **fp\_enable** subroutine or **fp\_enable\_all** subroutine must be called to enable the appropriate traps.

The header file **fptrap.h** defines the following names for the individual bits in the floating-point trap control:

<b>TRP_INVALID</b>	Invalid Operation Summary
<b>TRP_DIV_BY_ZERO</b>	Divide by Zero
<b>TRP_OVERFLOW</b>	Overflow
<b>TRP_UNDERFLOW</b>	Underflow
<b>TRP_INEXACT</b>	Inexact Result

### Parameters

*Mask*                      A 32-bit pattern that identifies floating-point traps.

### Return Values

The **fp\_any\_enable** subroutine returns 1 if any floating-point traps are enabled. Otherwise, 0 is returned.

The **fp\_is\_enabled** subroutine returns 1 if the floating-point traps specified by the *Mask* parameter are enabled. Otherwise, 0 is returned.

The **fp\_enable\_all** subroutine enables all floating-point traps.

The **fp\_enable** subroutine enables all floating-point traps specified by the *Mask* parameter.

The **fp\_disable\_all** subroutine disables all floating-point traps.

The **fp\_disable** subroutine disables all floating-point traps specified by the *Mask* parameter.

## Related Information

The **fp\_clr\_flag**, **fp\_set\_flag**, **fp\_read\_flag**, **fp\_swap\_flag** (“fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine”) subroutine, **fp\_invalid\_op**, **fp\_divbyzero**, **fp\_overflow**, **fp\_underflow**, **fp\_inexact**, **fp\_any\_xcp** (“fp\_invalid\_op, fp\_divbyzero, fp\_overflow, fp\_underflow, fp\_inexact, fp\_any\_xcp Subroutine” on page 299) subroutines, **fp\_iop\_snan**, **fp\_iop\_infsinf**, **fp\_iop\_infdfinf**, **fp\_iop\_zrdzr**, **fp\_iop\_infmzr**, **fp\_iop\_invcmp** (“fp\_iop\_snan, fp\_iop\_infsinf, fp\_iop\_infdfinf, fp\_iop\_zrdzr, fp\_iop\_infmzr, fp\_iop\_invcmp, fp\_iop\_sqrt, fp\_iop\_convert, or fp\_iop\_vxsoft Subroutines” on page 300) subroutines, **fp\_read\_rnd**, and **fp\_swap\_rnd** (“fp\_read\_rnd or fp\_swap\_rnd Subroutine” on page 302) subroutines, **fp\_trap** (“fp\_trap Subroutine” on page 306) subroutine.

Floating-Point Processor in *Assembler Language Reference*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine

### Purpose

Allows operations on the floating-point exception flags.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <float.h>
#include <fpxcp.h>
```

```
void fp_clr_flag( Mask)
fpflag_t Mask;
void fp_set_flag(Mask)
fpflag_t Mask;
fpflag_t fp_read_flag( )
fpflag_t fp_swap_flag(Mask)
fpflag_t Mask;
```

### Description

These subroutines aid in determining both when an exception has occurred and the exception type. These subroutines can be called explicitly around blocks of code that may cause a floating-point exception.

According to the *IEEE Standard for Binary Floating-Point Arithmetic*, the following types of floating-point operations must be signaled when detected in a floating-point operation:

- Invalid operation
- Division by zero
- Overflow
- Underflow
- Inexact

An invalid operation occurs when the result cannot be represented (for example, a **sqrt** operation on a number less than 0).

The *IEEE Standard for Binary Floating-Point Arithmetic* states: "For each type of exception, the implementation shall provide a status flag that shall be set on any occurrence of the corresponding exception when no corresponding trap occurs. It shall be reset only at the user's request. The user shall be able to test and to alter the status flags individually, and should further be able to save and restore all five at one time."

Floating-point operations can set flags in the floating-point exception status but cannot clear them. Users can clear a flag in the floating-point exception status using an explicit software action such as the **fp\_swap\_flag** (0) subroutine.

The **fpxcp.h** file defines the following names for the flags indicating floating-point exception status:

<b>FP_INVALID</b>	Invalid operation summary
<b>FP_OVERFLOW</b>	Overflow
<b>FP_UNDERFLOW</b>	Underflow
<b>FP_DIV_BY_ZERO</b>	Division by 0
<b>FP_INEXACT</b>	Inexact result

In addition to these flags, the operating system supports additional information about the cause of an invalid operation exception. The following flags also indicate floating-point exception status and defined in the **fpxcp.h** file. The flag number for each exception type varies, but the mnemonics are the same for all ports. The following invalid operation detail flags are not required for conformance to the IEEE floating-point exceptions standard:

<b>FP_INV_SNAN</b>	Signaling NaN
<b>FP_INV_ISI</b>	INF - INF
<b>FP_INV_IDI</b>	INF / INF
<b>FP_INV_ZDZ</b>	0 / 0
<b>FP_INV_IMZ</b>	INF x 0
<b>FP_INV_CMP</b>	Unordered compare
<b>FP_INV_SQRT</b>	Square root of a negative number
<b>FP_INV_CVI</b>	Conversion to integer error
<b>FP_INV_VXSOFT</b>	Software request

## Parameters

*Mask* A 32-bit pattern that identifies floating-point exception flags.

## Return Values

The **fp\_clr\_flag** subroutine resets the exception status flags defined by the *Mask* parameter to 0 (false). The remaining flags in the exception status are unchanged.

The **fp\_set\_flag** subroutine sets the exception status flags defined by the *Mask* parameter to 1 (true). The remaining flags in the exception status are unchanged.

The **fp\_read\_flag** subroutine returns the current floating-point exception status. The flags in the returned exception status can be tested using the flag definitions above. You can test individual flags or sets of flags.



The **fp\_swap\_flag** subroutine writes the *Mask* parameter into the floating-point status and returns the floating-point exception status from before the write.

Users set or reset multiple exception flags using **fp\_set\_flag** and **fp\_clr\_flag** by ANDing or ORing definitions for individual flags. For example, the following resets both the overflow and inexact flags:

```
fp_clr_flag (FP_OVERFLOW | FP_INEXACT)
```

## Related Information

The **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable**, or **fp\_disable\_all** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) subroutine, **fp\_any\_xcp**, **fp\_divbyzero**, **fp\_inexact**, **fp\_invalid\_op**, **fp\_overflow**, **fp\_underflow** (“fp\_invalid\_op, fp\_divbyzero, fp\_overflow, fp\_underflow, fp\_inexact, fp\_any\_xcp Subroutine” on page 299) subroutines, **fp\_iop\_infdinf**, **fp\_iop\_infmzr**, **fp\_iop\_infsinf**, **fp\_iop\_invcmp**, **fp\_iop\_snan**, or **fp\_iop\_zrdzr** (“fp\_iop\_snan, fp\_iop\_infsinf, fp\_iop\_infdinf, fp\_iop\_zrdzr, fp\_iop\_infmzr, fp\_iop\_invcmp, fp\_iop\_sqrt, fp\_iop\_convert, or fp\_iop\_vxsoft Subroutines” on page 300) subroutines, **fp\_read\_rnd** or **fp\_swap\_rnd** (“fp\_read\_rnd or fp\_swap\_rnd Subroutine” on page 302) subroutine.

Floating-Point Exceptions Overview and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fp\_cpusync Subroutine

### Purpose

Queries or changes the floating-point exception enable (FE) bit in the Machine Status register (MSR).

**Note:** This subroutine has been replaced by the **fp\_trapstate** (“fp\_trapstate Subroutine” on page 308) subroutine. The **fp\_cpusync** subroutine is supported for compatibility, but the **fp\_trapstate** subroutine should be used for development.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fptrap.h>
```

```
int fp_cpusync ( Flag );  
int Flag;
```

### Description

The **fp\_cpusync** subroutine is a service routine used to query, set, or reset the Machine Status Register (MSR) floating-point exception enable (FE) bit. The MSR FE bit determines whether a processor runs in pipeline or serial mode. Floating-point traps can only be generated by the hardware when the processor is in synchronous mode.

The **fp\_cpusync** subroutine changes only the MSR FE bit. It is a service routine for use in developing custom floating-point exception-handling software. If you are using the **fp\_enable** or **fp\_enable\_all** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) subroutine or the **fp\_sh\_trap\_info** or **fp\_sh\_set\_stat** (“fp\_sh\_info, fp\_sh\_trap\_info, or fp\_sh\_set\_stat Subroutine” on page 303) subroutine, you must use the **fp\_trap** (“fp\_trap Subroutine” on page 306) subroutine to place the process in serial mode.

## Parameters

<i>Flag</i>	Specifies to query or modify the MSR FE bit:
<b>FP_SYNC_OFF</b>	Sets the FE bit in the MSR to Off, which disables floating-point exception processing immediately.
<b>FP_SYNC_ON</b>	Sets the FE bit in the MSR to On, which enables floating-exception processing for the next floating-point operation.
<b>FP_SYNC_QUERY</b>	Returns the current state of the process (either <b>FP_SYNC_ON</b> or <b>FP_SYNC_OFF</b> ) without modifying it.

If called with any other value, the **fp\_cpusync** subroutine returns **FP\_SYNC\_ERROR**.

## Return Values

If called with the **FP\_SYNC\_OFF** or **FP\_SYNC\_ON** flag, the **fp\_cpusync** subroutine returns a value indicating which flag was in the previous state of the process.

If called with the **FP\_SYNC\_QUERY** flag, the **fp\_cpusync** subroutine returns a value indicating the current state of the process, either the **FP\_SYNC\_OFF** or **FP\_SYNC\_ON** flag.

## Error Codes

If the **fp\_cpusync** subroutine is called with an invalid parameter, the subroutine returns **FP\_SYNC\_ERROR**. No other errors are reported.

## Related Information

The **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** (“**fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** Subroutine” on page 294) subroutine, **fp\_clr\_flag**, **fpset\_flag**, **fp\_read\_flag**, or **fp\_swap\_flag** (“**fp\_clr\_flag**, **fp\_set\_flag**, **fp\_read\_flag**, or **fp\_swap\_flag** Subroutine” on page 295) subroutine, **sigaction**, **sigvec**, or **signal** subroutine.

Floating-Point Processor in *Assembler Language Reference*.

Floating-Point Exceptions in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fp\_flush\_imprecise Subroutine

### Purpose

Forces imprecise signal delivery.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fptrap.h>
void fp_flush_imprecise ()
```

## Description

The **fp\_flush\_imprecise** subroutine forces any imprecise interrupts to be reported. To ensure that no signals are lost when a program voluntarily exits, use this subroutine in combination with the **atexit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine.

## Example

The following example illustrates using the **atexit** subroutine to run the **fp\_flush\_imprecise** subroutine before a program exits:

```
#include <fptrap.h>
#include <stdlib.h>
#include <stdio.h>
    if (0!=atexit(fp_flush_imprecise))
        puts ("Failure in atexit(fp_flush_imprecise) ");
```

## Related Information

The **atexit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) subroutine, **fp\_clr\_flag**, **fp\_read\_flag**, **fp\_swap\_flag**, or **fpset\_flag** (“fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine” on page 295) subroutine, **fp\_cpusync** (“fp\_cpusync Subroutine” on page 297) subroutine, **fp\_trap** (“fp\_trap Subroutine” on page 306) subroutine, **sigaction** subroutine.

Floating-Point Exceptions in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## **fp\_invalid\_op, fp\_divbyzero, fp\_overflow, fp\_underflow, fp\_inexact, fp\_any\_xcp** Subroutine

### Purpose

Tests to see if a floating-point exception has occurred.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <float.h>
#include <fpxcp.h>
int
fp_invalid_op()
int fp_divbyzero()
int fp_overflow()
int fp_underflow()
int
fp_inexact()
int fp_any_xcp()
```

### Description

These subroutines aid in determining when an exception has occurred and the exception type. These subroutines can be called explicitly after blocks of code that may cause a floating-point exception.

## Return Values

The **fp\_invalid\_op** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set. Otherwise, a value of 0 is returned.

The **fp\_divbyzero** subroutine returns a value of 1 if a floating-point divide-by-zero exception status flag is set. Otherwise, a value of 0 is returned.

The **fp\_overflow** subroutine returns a value of 1 if a floating-point overflow exception status flag is set. Otherwise, a value of 0 is returned.

The **fp\_underflow** subroutine returns a value of 1 if a floating-point underflow exception status flag is set. Otherwise, a value of 0 is returned.

The **fp\_inexact** subroutine returns a value of 1 if a floating-point inexact exception status flag is set. Otherwise, a value of 0 is returned.

The **fp\_any\_xcp** subroutine returns a value of 1 if a floating-point invalid operation, divide-by-zero, overflow, underflow, or inexact exception status flag is set. Otherwise, a value of 0 is returned.

## Related Information

The **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable fp\_disable\_all**, or **fp\_disable** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) subroutine, **fp\_clr\_flag**, **fp\_read\_flag**, **fp\_set\_flag**, or **fp\_swap\_flag** (“fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine” on page 295) subroutine, **fp\_read\_rnd** or **fp\_swap\_rnd** (“fp\_read\_rnd or fp\_swap\_rnd Subroutine” on page 302) subroutine.

Floating-Point Processor in *Assembler Language Reference*.

Floating-Point Exceptions and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## **fp\_iop\_snan, fp\_iop\_infsinf, fp\_iop\_infdinf, fp\_iop\_zrdzr, fp\_iop\_infmzr, fp\_iop\_invcmp, fp\_iop\_sqrt, fp\_iop\_convert, or fp\_iop\_vxsoft** Subroutines

### Purpose

Tests to see if a floating-point exception has occurred.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <float.h>
#include <fp MCP.h>
int fp_iop_snan()
int fp_iop_infsinf()
int
fp_iop_infdinf()
int fp_iop_zrdzr()
int
fp_iop_infmzr()
int fp_iop_invcmp()
```

```
int
fp_iop_sqrt()
int fp_iop_convert()
int
fp_iop_vxsoft ();
```

## Description

These subroutines aid in determining when an exception has occurred and the exception type. These subroutines can be called explicitly after blocks of code that may cause a floating-point exception.

## Return Values

The **fp\_iop\_snan** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to a signaling NaN (NaNs) flag. Otherwise, a value of 0 is returned.

The **fp\_iop\_infsinf** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to an INF-INF flag. Otherwise, a value of 0 is returned.

The **fp\_iop\_infdir** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to an INF/INF flag. Otherwise, a value of 0 is returned.

The **fp\_iop\_zrdzr** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to a 0.0/0.0 flag. Otherwise, a value of 0 is returned.

The **fp\_iop\_infmzr** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to an INF\*0.0 flag. Otherwise, a value of 0 is returned.

The **fp\_iop\_invcmp** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to a compare involving a NaN. Otherwise, a value of 0 is returned.

The **fp\_iop\_sqrt** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to the calculation of a square root of a negative number. Otherwise, a value of 0 is returned.

The **fp\_iop\_convert** subroutine returns a value of 1 if a floating-point invalid-operation exception status flag is set due to the conversion of a floating-point number to an integer, where the floating-point number was a NaN, an INF, or was outside the range of the integer. Otherwise, a value of 0 is returned.

The **fp\_iop\_vxsoft** subroutine returns a value of 1 if the VXSOFT detail bit is on. Otherwise, a value of 0 is returned.

---

## fp\_raise\_xcp Subroutine

### Purpose

Generates a floating-point exception.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fp_xcp.h>
```

```
int fp_raise_xcp( mask)
fpflag_t mask;
```

## Description

The **fp\_raise\_xcp** subroutine causes any floating-point exceptions defined by the *mask* parameter to be raised immediately. If the exceptions defined by the *mask* parameter are enabled and the program is running in serial mode, the signal for floating-point exceptions, **SIGFPE**, is raised.

If more than one exception is included in the *mask* variable, the exceptions are raised in the following order:

1. Invalid
2. Dividebyzero
3. Underflow
4. Overflow
5. Inexact

Thus, if the user exception handler does not disable further exceptions, one call to the **fp\_raise\_xcp** subroutine can cause the exception handler to be entered many times.

## Parameters

*mask* Specifies a 32-bit pattern that identifies floating-point traps.

## Return Values

The **fp\_raise\_xcp** subroutine returns 0 for normal completion and returns a nonzero value if an error occurs.

## Related Information

The **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** (“**fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** Subroutine” on page 294) subroutine, **fp\_clr\_flag**, **fp\_read\_flag**, **fp\_swap\_flag**, or **fpset\_flag** (“**fp\_clr\_flag**, **fp\_set\_flag**, **fp\_read\_flag**, or **fp\_swap\_flag** Subroutine” on page 295) subroutine, **fp\_cpusync** (“**fp\_cpusync** Subroutine” on page 297) subroutine, **fp\_trap** (“**fp\_trap** Subroutine” on page 306) subroutine, **sigaction** subroutine.

---

## fp\_read\_rnd or fp\_swap\_rnd Subroutine

### Purpose

Read and set the IEEE floating-point rounding mode.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <float.h>
```

```
fprnd_t fp_read_rnd()  
fprnd_t fp_swap_rnd( RoundMode)  
fprnd_t RoundMode;
```

### Description

The **fp\_read\_rnd** subroutine returns the current rounding mode. The **fp\_swap\_rnd** subroutine changes the rounding mode to the *RoundMode* parameter and returns the value of the rounding mode before the change.

Floating-point rounding occurs when the infinitely precise result of a floating-point operation cannot be represented exactly in the destination floating-point format (such as double-precision format).

The *IEEE Standard for Binary Floating-Point Arithmetic* allows floating-point numbers to be rounded in four different ways: round toward zero, round to nearest, round toward +INF, and round toward -INF. Once a rounding mode is selected it affects all subsequent floating-point operations until another rounding mode is selected.

**Note:** The default floating-point rounding mode is round to nearest. All C main programs begin with the rounding mode set to round to nearest.

The encodings of the rounding modes are those defined in the *ANSI C Standard*. The **float.h** file contains definitions for the rounding modes. Below is the **float.h** definition, the *ANSI C Standard* value, and a description of each rounding mode.

float.h Definition	ANSI Value	Description
FP_RND_RZ	0	Round toward 0
FP_RND_RN	1	Round to nearest
FP_RND_RP	2	Round toward +INF
FP_RND_RM	3	Round toward -INF

The **fp\_swap\_rnd** subroutine can be used to swap rounding modes by saving the return value from **fp\_swap\_rnd(RoundMode)**. This can be useful in functions that need to force a specific rounding mode for use during the function but wish to restore the caller's rounding mode on exit. Below is a code fragment that accomplishes this action:

```
save_mode = fp_swap_rnd (new_mode);
...desired code using new_mode
(void) fp_swap_rnd(save_mode); /*restore caller's mode*/
```

## Parameters

*RoundMode* Specifies one of the following modes: **FP\_RND\_RZ**, **FP\_RND\_RN**, **FP\_RND\_RP**, or **FP\_RND\_RM**.

## Related Information

The **floor**, **ceil**, **nearest**, **trunc**, **rint**, **itrunc**, **uitrunc**, **fmod**, or **fabs** (“floor, floorf, floorl, floord32, floord64, floord128, nearest, trunc, itrunc, and uitrunc Subroutines” on page 277) subroutine, **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) subroutine, **fp\_clr\_flag**, **fp\_read\_flag**, **fp\_set\_flag**, or **fp\_swap\_flag** (“fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine” on page 295) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fp\_sh\_info, fp\_sh\_trap\_info, or fp\_sh\_set\_stat Subroutine

### Purpose

From within a floating-point signal handler, determines any floating-point exception that caused the trap in the process and changes the state of the Floating-Point Status and Control register (FPSCR) in the user process.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <fp MCP.h>
#include <fp trap.h>
#include <signal.h>
```

```
void fp_sh_info( scp, fcp, struct_size)
struct sigcontext *scp;
struct fp_sh_info *fcp;
size_t struct_size;
```

```
void fp_sh_trap_info( scp, fcp)
struct sigcontext *scp;
struct fp_ctx *fcp;
```

```
void fp_sh_set_stat( scp, fpscr)
struct sigcontext *scp;
fpstat_t fpscr;
```

## Description

These subroutines are for use within a user-written signal handler. They return information about the process that was running at the time the signal occurred, and they update the Floating-Point Status and Control register for the process.

**Note:** The **fp\_sh\_trap\_info** subroutine is maintained for compatibility only. It has been replaced by the **fp\_sh\_info** subroutine, which should be used for development.

These subroutines operate only on the state of the user process that was running at the time the signal was delivered. They read and write the **sigcontext** structure. They do not change the state of the signal handler process itself.

The state of the signal handler process can be modified by the **fp\_any\_enable**, **fp\_is\_enabled**, **fp\_enable\_all**, **fp\_enable**, **fp\_disable\_all**, or **fp\_disable** subroutine.

### fp\_sh\_info

The **fp\_sh\_info** subroutine returns information about the process that caused the trap by means of a floating-point context (**fp\_sh\_info**) structure. This structure contains the following information:

```
typedef struct fp_sh_info {
fpstat_t      fpscr;
fpflag_t      trap;
short         trap_mode;
char          flags;
char          extra;
} fp_sh_info_t;
```

The fields are:

fpscr	The Floating-Point Status and Control register (FPSCR) in the user process at the time the interrupt occurred.
trap	A mask indicating the trap or traps that caused the signal handler to be entered. This mask is the logical OR operator of the enabled floating-point exceptions that occurred to cause the trap. This mask can have up to two exceptions; if there are two, the <b>INEXACT</b> signal must be one of them. If the mask is 0, the <b>SIGFPE</b> signal was raised not by a floating-point operation, but by the <b>kill</b> or <b>raise</b> subroutine or the <b>kill</b> command.



trap_mode	The trap mode in effect in the process at the time the signal handler was entered. The values returned in the <b>fp_sh_info.trap_mode</b> file use the following argument definitions:  <b>FP_TRAP_OFF</b> Trapping off  <b>FP_TRAP_SYNC</b> Precise trapping on  <b>FP_TRAP_IMP_REC</b> Recoverable imprecise trapping on  <b>FP_TRAP_IMP</b> Non-recoverable imprecise trapping on
flags	This field is interpreted as an array of bits and should be accessed with masks. The following mask is defined:  <b>FP_IAR_STAT</b> If the value of the bit at this mask is 1, the exception was precise and the IAR points to the instruction that caused the exception. If the value bit at this mask is 0, the exception was imprecise.

## fp\_sh\_trap\_info

The **fp\_sh\_trap\_info** subroutine is maintained for compatibility only. The **fp\_sh\_trap\_info** subroutine returns information about the process that caused the trap by means of a floating-point context (**fp\_ctx**) structure. This structure contains the following information:

```
fpstat_t fpscr;
fpflag_t trap;
```

The fields are:

fpscr	The Floating-Point Status and Control register (FPSCR) in the user process at the time the interrupt occurred.
trap	A mask indicating the trap or traps that caused the signal handler to be entered. This mask is the logical OR operator of the enabled floating-point exceptions that occurred to cause the trap. This mask can have up to two exceptions; if there are two, the <b>INEXACT</b> signal must be one of them. If the mask is 0, the <b>SIGFPE</b> signal was raised not by a floating-point operation, but by the <b>kill</b> or <b>raise</b> subroutine or the <b>kill</b> command.

## fp\_sh\_set\_stat

The **fp\_sh\_set\_stat** subroutine updates the Floating-Point Status and Control register (FPSCR) in the user process with the value in the **fpscr** field.

The signal handler must either clear the exception bit that caused the trap to occur or disable the trap to prevent a recurrence. If the instruction generated more than one exception, and the signal handler clears only one of these exceptions, a signal is raised for the remaining exception when the next floating-point instruction is executed in the user process.

## Parameters

<i>fcpx</i>	Specifies a floating-point context structure.
<i>scpx</i>	Specifies a <b>sigcontext</b> structure for the interrupt.
<i>struct_size</i>	Specifies the size of the <b>fp_sh_info</b> structure.
<i>fpscr</i>	Specifies which Floating-Point Status and Control register to update.

## Related Information

The **fp\_any\_enable**, **fp\_disable\_all**, **fp\_disable**, **fp\_enable\_all**, **fp\_enable**, or **fp\_is\_enabled** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) subroutine, **fp\_clr\_flag**, **fp\_read\_flag**, **fp\_set\_flag**, or **fp\_swap\_flag** (“fp\_clr\_flag, fp\_set\_flag, fp\_read\_flag, or fp\_swap\_flag Subroutine” on page 295) subroutine, **fp\_trap** (“fp\_trap Subroutine”) subroutine.

Floating-Point Exceptions in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fp\_trap Subroutine

### Purpose

Queries or changes the mode of the user process to allow floating-point exceptions to generate traps.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fptrap.h>
```

```
int fp_trap( flag)
int flag;
```

### Description

The **fp\_trap** subroutine queries and changes the mode of the user process to allow or disallow floating-point exception trapping. Floating-point traps can only be generated when a process is executing in a traps-enabled mode.

The default state is to execute in pipelined mode and not to generate floating-point traps.

**Note:** The **fp\_trap** routines only change the execution state of the process. To generate floating-point traps, you must also enable traps. Use the **fp\_enable** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) and **fp\_enable\_all** subroutines to enable traps.

Before calling the **fp\_trap(FP\_TRAP\_SYNC)** routine, previous floating-point operations can set to True certain exception bits in the Floating-Point Status and Control register (FPSCR). Enabling these Cexceptions and calling the **fp\_trap(FP\_TRAP\_SYNC)** routine does not cause an immediate trap to occur. That is, the operation of these traps is edge-sensitive, not level-sensitive.

The **fp\_trap** subroutine does not clear the exception history. You can query this history by using any of the following subroutines:

- **fp\_any\_xcp**
- **fp\_divbyzero**
- **fp\_iop\_convert**
- **fp\_iop\_infdinf**
- **fp\_iop\_infmzr**
- **fp\_iop\_infsinf**
- **fp\_iop\_invcmp**
- **fp\_iop\_snan**

- **fp\_iop\_sqrt**
- **fp\_iop\_vxsoft**
- **fp\_iop\_zrdzr**
- **fp\_inexact**
- **fp\_invalid\_op**
- **fp\_overflow**
- **fp\_underflow**

## Parameters

*flag* Specifies a query of or change in the mode of the user process:

### **FP\_TRAP\_OFF**

Puts the user process into trapping-off mode and returns the previous mode of the process, either **FP\_TRAP\_SYNC**, **FP\_TRAP\_IMP**, **FP\_TRAP\_IMP\_REC**, or **FP\_TRAP\_OFF**.

### **FP\_TRAP\_QUERY**

Returns the current mode of the user process.

### **FP\_TRAP\_SYNC**

Puts the user process into precise trapping mode and returns the previous mode of the process.

### **FP\_TRAP\_IMP**

Puts the user process into non-recoverable imprecise trapping mode and returns the previous mode.

### **FP\_TRAP\_IMP\_REC**

Puts the user process into recoverable imprecise trapping mode and returns the previous mode.

### **FP\_TRAP\_FASTMODE**

Puts the user process into the fastest trapping mode available on the hardware platform.

**Note:** Some hardware models do not support all modes. If an unsupported mode is requested, the **fp\_trap** subroutine returns **FP\_TRAP\_UNIMPL**.

## Return Values

If called with the **FP\_TRAP\_OFF**, **FP\_TRAP\_IMP**, **FP\_TRAP\_IMP\_REC**, or **FP\_TRAP\_SYNC** flag, the **fp\_trap** subroutine returns a value indicating which flag was in the previous mode of the process if the hardware supports the requested mode. If the hardware does not support the requested mode, the **fp\_trap** subroutine returns **FP\_TRAP\_UNIMPL**.

If called with the **FP\_TRAP\_QUERY** flag, the **fp\_trap** subroutine returns a value indicating the current mode of the process, either the **FP\_TRAP\_OFF**, **FP\_TRAP\_IMP**, **FP\_TRAP\_IMP\_REC**, or **FP\_TRAP\_SYNC** flag.

If called with **FP\_TRAP\_FASTMODE**, the **fp\_trap** subroutine sets the fastest mode available and returns the mode selected.

## Error Codes

If the **fp\_trap** subroutine is called with an invalid parameter, the subroutine returns **FP\_TRAP\_ERROR**.

If the requested mode is not supported on the hardware platform, the subroutine returns **FP\_TRAP\_UNIMPL**.

---

## fp\_trapstate Subroutine

### Purpose

Queries or changes the trapping mode in the Machine Status register (MSR).

**Note:** This subroutine replaces the **fp\_cpusync** (“fp\_cpusync Subroutine” on page 297) subroutine. The **fp\_cpusync** subroutine is supported for compatibility, but the **fp\_trapstate** subroutine should be used for development.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fptrap.h>
int fp_trapstate (int)
```

### Description

The **fp\_trapstate** subroutine is a service routine used to query or set the trapping mode. The trapping mode determines whether floating-point exceptions can generate traps, and can affect execution speed. See Floating-Point Exceptions Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for a description of precise and imprecise trapping modes. Floating-point traps can be generated by the hardware only when the processor is in a traps-enabled mode.

The **fp\_trapstate** subroutine changes only the trapping mode. It is a service routine for use in developing custom floating-point exception-handling software. If you are using the **fp\_enable** (“fp\_any\_enable, fp\_is\_enabled, fp\_enable\_all, fp\_enable, fp\_disable\_all, or fp\_disable Subroutine” on page 294) or **fp\_enable\_all** subroutine or the **fp\_sh\_info** (“fp\_sh\_info, fp\_sh\_trap\_info, or fp\_sh\_set\_stat Subroutine” on page 303) or **fp\_sh\_set\_stat** subroutine, you must use the **fp\_trap** (“fp\_trap Subroutine” on page 306) subroutine to change the process’ trapping mode.

### Parameters

*flag* Specifies a query of, or change in, the trap mode:

**FP\_TRAPSTATE\_OFF**

Sets the trapping mode to Off and returns the previous mode.

**FP\_TRAPSTATE\_QUERY**

Returns the current trapping mode without modifying it.

**FP\_TRAPSTATE\_IMP**

Puts the process in non-recoverable imprecise trapping mode and returns the previous state.

**FP\_TRAPSTATE\_IMP\_REC**

Puts the process in recoverable imprecise trapping mode and returns the previous state.

**FP\_TRAPSTATE\_PRECISE**

Puts the process in precise trapping mode and returns the previous state.

**FP\_TRAPSTATE\_FASTMODE**

Puts the process in the fastest trap-generating mode available on the hardware platform and returns the state selected.

**Note:** Some hardware models do not support all modes. If an unsupported mode is requested, the **fp\_trapstate** subroutine returns **FP\_TRAP\_UNIMPL** and the trapping mode is not changed.

## Return Values

If called with the `FP_TRAPSTATE_OFF`, `FP_TRAPSTATE_IMP`, `FP_TRAPSTATE_IMP_REC`, or `FP_TRAPSTATE_PRECISE` flag, the `fp_trapstate` subroutine returns a value indicating the previous mode of the process. The value may be `FP_TRAPSTATE_OFF`, `FP_TRAPSTATE_IMP`, `FP_TRAPSTATE_IMP_REC`, or `FP_TRAPSTATE_PRECISE`. If the hardware does not support the requested mode, the `fp_trapstate` subroutine returns `FP_TRAP_UNIMPL`.

If called with the `FP_TRAP_QUERY` flag, the `fp_trapstate` subroutine returns a value indicating the current mode of the process. The value may be `FP_TRAPSTATE_OFF`, `FP_TRAPSTATE_IMP`, `FP_TRAPSTATE_IMP_REC`, or `FP_TRAPSTATE_PRECISE`.

If called with the `FP_TRAPSTATE_FASTMODE` flag, the `fp_trapstate` subroutine returns a value indicating which mode was selected. The value may be `FP_TRAPSTATE_OFF`, `FP_TRAPSTATE_IMP`, `FP_TRAPSTATE_IMP_REC`, or `FP_TRAPSTATE_PRECISE`.

## Related Information

The `fp_any_enable`, `fp_disable_all`, `fp_disable`, `fp_enable_all`, `fp_enable`, or `fp_is_enabled` (“`fp_any_enable`, `fp_is_enabled`, `fp_enable_all`, `fp_enable`, `fp_disable_all`, or `fp_disable` Subroutine” on page 294) subroutine, `fp_clr_flag`, `fp_read_flag`, `fpset_flag`, or `fp_swap_flag` (“`fp_clr_flag`, `fp_set_flag`, `fp_read_flag`, or `fp_swap_flag` Subroutine” on page 295) subroutine, `sigaction`, `signal`, or `sigvec` subroutine.

The Floating-Point Processor in *Assembler Language Reference*.

Floating-Point Exceptions in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fpclassify Macro

### Purpose

Classifies real floating type.

### Syntax

```
#include <math.h>
```

```
int fpclassify(x)
real-floating x;
```

### Description

The `fpclassify` macro classifies the `x` parameter as NaN, infinite, normal, subnormal, zero, or into another implementation-defined category. An argument represented in a format wider than its semantic type is converted to its semantic type. Classification is based on the type of the argument.

### Parameters

`x` Specifies the value to be classified.

## Return Values

The `fpclassify` macro returns the value of the number classification macro appropriate to the value of its argument.

## Related Information

“isfinite Macro” on page 566, “isinf Subroutine” on page 568, “class, \_class, finite, isnan, or unordered Subroutines” on page 167, “isnormal Macro” on page 570.

The signbit Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fread or fwrite Subroutine

### Purpose

Reads and writes binary files.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
size_t fread ( (void *) Pointer, Size, NumberOfItems, Stream ("Parameters" on page 311))
size_t Size, NumberOfItems ("Parameters" on page 311);
FILE *Stream ("Parameters" on page 311);
size_t fwrite (Pointer, Size, NumberOfItems, Stream ("Parameters" on page 311))
const void *Pointer ("Parameters" on page 311);
size_t Size, NumberOfItems ("Parameters" on page 311);
FILE *Stream ("Parameters" on page 311);
```

### Description

The **fread** subroutine copies the number of data items specified by the *NumberOfItems* parameter from the input stream into an array beginning at the location pointed to by the *Pointer* parameter. Each data item has the form *\*Pointer*.

The **fread** subroutine stops copying bytes if an end-of-file (EOF) or error condition is encountered while reading from the input specified by the *Stream* parameter, or when the number of data items specified by the *NumberOfItems* parameter have been copied. This subroutine leaves the file pointer of the *Stream* parameter, if defined, pointing to the byte following the last byte read. The **fread** subroutine does not change the contents of the *Stream* parameter.

The *st\_atime* field will be marked for update by the first successful run of the **fgetc** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **fgets** (“gets or fgets Subroutine” on page 434), **fgetwc** (“getwc, fgetwc, or getwchar Subroutine” on page 477), **fgetws** (“getws or fgetws Subroutine” on page 480), **fread**, **fscanf**, **getc** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **getchar** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **gets** (“gets or fgets Subroutine” on page 434), or **scanf** subroutine using a stream that returns data not supplied by a prior call to the **ungetc** or **ungetwc** subroutine.

**Note:** The **fread** subroutine is a buffered **read** subroutine library call. It reads data in 4KB blocks. For tape block sizes greater than 4KB, use the **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine and **read** subroutine.

The **fwrite** subroutine writes items from the array pointed to by the *Pointer* parameter to the stream pointed to by the *Stream* parameter. Each item’s size is specified by the *Size* parameter. The **fwrite** subroutine writes the number of items specified by the *NumberOfItems* parameter. The file-position

indicator for the stream is advanced by the number of bytes successfully written. If an error occurs, the resulting value of the file-position indicator for the stream is indeterminate.

The **fwrite** subroutine appends items to the output stream from the array pointed to by the *Pointer* parameter. The **fwrite** subroutine appends as many items as specified in the *NumberOfItems* parameter.

The **fwrite** subroutine stops writing bytes if an error condition is encountered on the stream, or when the number of items of data specified by the *NumberOfItems* parameter have been written. The **fwrite** subroutine does not change the contents of the array pointed to by the *Pointer* parameter.

The *st\_ctime* and *st\_mtime* fields will be marked for update between the successful run of the **fwrite** subroutine and the next completion of a call to the **fflush** (“fclose or fflush Subroutine” on page 253) or **fclose** subroutine on the same stream, the next call to the **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, or the next call to the **abort** (“abort Subroutine” on page 2) subroutine.

## Parameters

<i>Pointer</i>	Points to an array.
<i>Size</i>	Specifies the size of the variable type of the array pointed to by the <i>Pointer</i> parameter. The <i>Size</i> parameter can be considered the same as a call to <b>sizeof</b> subroutine.
<i>NumberOfItems</i>	Specifies the number of items of data.
<i>Stream</i>	Specifies the input or output stream.

## Return Values

The **fread** and **fwrite** subroutines return the number of items actually transferred. If the *NumberOfItems* parameter contains a 0, no characters are transferred, and a value of 0 is returned. If the *NumberOfItems* parameter contains a negative number, it is translated to a positive number, since the *NumberOfItems* parameter is of the unsigned type.

## Error Codes

If the **fread** subroutine is unsuccessful because the I/O stream is unbuffered or data needs to be read into the I/O stream’s buffer, it returns one or more of the following error codes:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> flag is set for the file descriptor specified by the <i>Stream</i> parameter, and the process would be delayed in the <b>fread</b> operation.
<b>EBADF</b>	Indicates that the file descriptor specified by the <i>Stream</i> parameter is not a valid file descriptor open for reading.
<b>EINTR</b>	Indicates that the read operation was terminated due to receipt of a signal, and no data was transferred.

**Note:** Depending upon which library routine the application binds to, this subroutine may return **EINTR**. Refer to the **signal** subroutine regarding **sa\_restart**.

<b>EIO</b>	Indicates that the process is a member of a background process group attempting to perform a read from its controlling terminal, and either the process is ignoring or blocking the <b>SIGTTIN</b> signal or the process group has no parent process.
<b>ENOMEM</b>	Indicates that insufficient storage space is available.
<b>ENXIO</b>	Indicates that a request was made of a nonexistent device.

If the **fwrite** subroutine is unsuccessful because the I/O stream is unbuffered or the I/O stream’s buffer needs to be flushed, it returns one or more of the following error codes:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> or <b>O_NDELAY</b> flag is set for the file descriptor specified by the <i>Stream</i> parameter, and the process is delayed in the write operation.
<b>EBADF</b>	Indicates that the file descriptor specified by the <i>Stream</i> parameter is not a valid file descriptor open for writing.
<b>EFBIG</b>	Indicates that an attempt was made to write a file that exceeds the file size of the process limit or the systemwide maximum file size.
<b>EINTR</b>	Indicates that the write operation was terminated due to the receipt of a signal, and no data was transferred.
<b>EIO</b>	Indicates that the process is a member of a background process group attempting to perform a write to its controlling terminal, the <b>TOSTOP</b> signal is set, the process is neither ignoring nor blocking the <b>SIGTTOU</b> signal, and the process group of the process is orphaned.
<b>ENOSPC</b>	Indicates that there was no free space remaining on the device containing the file.
<b>EPIPE</b>	Indicates that an attempt is made to write to a pipe or first-in-first-out (FIFO) process that is not open for reading by any process. A <b>SIGPIPE</b> signal is sent to the process.

The **fwrite** subroutine is also unsuccessful due to the following error conditions:

<b>ENOMEM</b>	Indicates that insufficient storage space is available.
<b>ENXIO</b>	Indicates that a request was made of a nonexistent device, or the request was outside the capabilities of the device.

## Related Information

The **abort** (“abort Subroutine” on page 2) subroutine, **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fflush** or **fclose** (“fclose or fflush Subroutine” on page 253) subroutine, **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **getc**, **getchar**, **fgetc**, or **getw** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **getwc**, **fgetwc**, or **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **gets** or **fgets** (“gets or fgets Subroutine” on page 434) subroutine, **getws** or **fgetws** (“getws or fgetws Subroutine” on page 480) subroutine, **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **print**, **fprintf**, or **sprintf** (“printf, fprintf, sprintf, snprintf, vsprintf, vprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **putc**, **putchar**, **fputc**, or **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **putwc**, **putwchar**, or **fputwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **puts** or **fputs** (“puts or fputs Subroutine” on page 1331) subroutine, **putws** or **fputws** (“putws or fputws Subroutine” on page 1341) subroutine, **read** subroutine, **scanf**, **fscanf**, **sscanf**, or **wscanf** subroutine, **ungetc** or **ungetwc** subroutine, **write** subroutine.

The Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## freehostent Subroutine

### Purpose

To free memory allocated by `getipnodebyname` and `getipnodebyaddr`.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void freehostent (ptr)
struct hostent * ptr;
```



## Description

The **freehostent** subroutine frees any dynamic storage pointed to by elements of *ptr*. This includes the **hostent** structure and the data areas pointed to by the *h\_name*, *h\_addr\_list*, and *h\_aliases* members of the **hostent** structure.

## Related Information

The **getipnodebyaddr** subroutine and **getipnodebyname** subroutine.

---

## freelmb Subroutine

### Purpose

Returns a block of memory allocated by **alloclmb()** to the system.

### Syntax

```
#include <sys/dr.h>
```

```
int freelmb(long long laddr)
```

### Description

The **freelmb()** subroutine returns a block of memory, allocated by **alloclmb()**, for general system use.

### Parameters

*laddr*                    A previously allocated LMB address.

### Execution Environment

This **freelmb()** interface should only be called from the process environment.

### Return Values

0                                    The LMB is successfully freed.

### Error Codes

<b>ENOTSUP</b>	LMB allocation not supported on this system.
<b>EINVAL</b>	<i>laddr</i> does not describe a previously allocated LMB.
<b>EINVAL</b>	Not in the process environment.

### Related Information

“alloclmb Subroutine” on page 68

---

## frevoke Subroutine

### Purpose

Revokes access to a file by other processes.

### Library

Standard C Library (**libc.a**)

## Syntax

```
int frevoke ( FileDescriptor)
int FileDescriptor;
```

## Description

The **frevoke** subroutine revokes access to a file by other processes.

All accesses to the file are revoked, except through the file descriptor specified by the *FileDescriptor* parameter to the **frevoke** subroutine. Subsequent attempts to access the file, using another file descriptor established before the **frevoke** subroutine was called, fail and cause the process to receive a return value of -1, and the **errno** global variable is set to **EBADF** .

A process can revoke access to a file only if its effective user ID is the same as the file owner ID or if the invoker has root user authority.

**Note:** The **frevoke** subroutine has no affect on subsequent attempts to open the file. To ensure exclusive access to the file, the caller should change the mode of the file before issuing the **frevoke** subroutine. Currently the **frevoke** subroutine works only on terminal devices.

## Parameters

*FileDescriptor*                      A file descriptor returned by a successful **open** subroutine.

## Return Values

Upon successful completion, the **frevoke** subroutine returns a value of 0.

If the **frevoke** subroutine fails, it returns a value of -1 and the **errno** global variable is set to indicate the error.

## Error Codes

The **frevoke** subroutine fails if the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> value is not the valid file descriptor of a terminal.
<b>EPERM</b>	The effective user ID of the calling process is not the same as the file owner ID.
<b>EINVAL</b>	Revocation of access rights is not implemented for this file.

---

## frexpd32, frexpd64, and frexpd128 Subroutines

### Purpose

Extracts the mantissa and exponent from a decimal floating-point number.

### Syntax

```
#include <math.h>

_Decimal32 frexpd32 (num, exp)
_Decimal32 num;
int *exp;

_Decimal64 frexpd64 (num, exp)
_Decimal64 num;
int *exp;
```

```
_Decimal128 frexpd128 (num, exp)
_Decimal128 num;
int *exp;
```

## Description

The **frexpd32**, **frexpd64**, and **frexpd128** subroutines break a decimal floating-point number into a mantissa and an integral power of 10. The integer exponent is stored in the **int** object pointed to by the **exp** parameter.

## Parameters

<b>num</b>	Specifies the decimal floating-point number to be broken into a mantissa and an integral power of 10.
<b>exp</b>	Points to where the integer exponent is stored.

## Return Values

For finite arguments, the **frexpd32**, **frexpd64**, and **frexpd128** subroutines return the mantissa value in the **x** parameter, such that the **num** parameter equals the **x** parameter times 10 raised to the power **exp** parameter.

If *num* is NaN, a NaN is returned, and the value of the *\*exp* is not specified.

If *num* is  $\pm 0$ ,  $\pm 0$  is returned, and the value of the *\*exp* is 0.

If *num* is  $\pm \text{Inf}$ , *num* is returned, and the value of the *\*exp* is not specified.

## Related Information

The “modf, modff, modfl, modfd32, modfd64, and modfd128 Subroutines” on page 848

---

## frexpf, frexpl, or frexp Subroutine

### Purpose

Extracts the mantissa and exponent from a double precision number.

### Syntax

```
#include <math.h>
```

```
float frexpf (num, exp)
float num;
int *exp;
```

```
long double frexpl (num, exp)
long double num;
int *exp;
```

```
double frexp (num, exp)
double num;
int *exp;
```

### Description

The **frexpf**, **frexpl**, and **frexp** subroutines break a floating-point number *num* into a normalized fraction and an integral power of 2. The integer exponent is stored in the **int** object pointed to by *exp*.

## Parameters

*num* Specifies the floating-point number to be broken into a normalized fraction and an integral power of 2.  
*exp* Points to where the integer exponent is stored.

## Return Values

For finite arguments, the **frexpf**, **frexpl**, and **frexp** subroutines return the value *x*, such that *x* has a magnitude in the interval  $[\frac{1}{2}, 1)$  or 0, and *num* equals *x* times 2 raised to the power *exp*.

If *num* is NaN, a NaN is returned, and the value of *\*exp* is unspecified.

If *num* is  $\pm 0$ ,  $\pm 0$  is returned, and the value of *\*exp* is 0.

If *num* is  $\pm \text{Inf}$ , *num* is returned, and the value of *\*exp* is unspecified.

## Related Information

“class, \_class, finite, isnan, or unordered Subroutines” on page 167 and “modf, modff, modfl, modfd32, modfd64, and modfd128 Subroutines” on page 848

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## fscntl Subroutine

### Purpose

Controls file system control operations.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <j2/j2_cntl.h>
#include <sys/vmount.h>
```

```
int fscntl ( vfs_id, Command, Argument, ArgumentSize)
int vfs_id;
int Command;
char *Argument;
int ArgumentSize;
```

### Description

The **fscntl** subroutine performs a variety of file system-specific functions. These functions typically require root user authority.

The Enhanced Journaled File System (JFS2) supports several *Command* values that can be used by applications. Each of these *Command* values requires root authority.

#### FSCNTL\_FREEZE

The file system specified by *vfs\_id* is “frozen” for a specified amount of time. The act of freezing a file system produces a nearly consistent on-disk image of the file system, and writes all dirty file system metadata and user data to the disk. In its frozen state, the file system is read-only, and anything that attempts to modify the file system or its contents must wait for the freeze to end. The

*Argument* is treated as an integral timeout value in seconds (instead of a pointer). The file system is thawed by **FSCNTL\_THAW** or when the timeout expires. The timeout, which must be a positive value, can be renewed using **FSCNTL\_REFREEZE**. The *ArgumentSize* must be 0.

**Note:** For all applications using this interface, use **FSCNTL\_THAW** to thaw the file system rather than waiting for the timeout to expire. If the timeout expires, an error log entry is generated as an advisory.

### **FSCNTL\_REFREEZE**

The file system specified by *vfs\_id*, which must be already frozen, has its timeout value reset. If the command is used on a file system that is not frozen, an error is returned. The *Argument* is treated as an integral timeout value in seconds (instead of a pointer). The file system is thawed by **FSCNTL\_THAW** or when the new timeout expires. The timeout must be a positive value. The *ArgumentSize* must be 0.

### **FSCNTL\_THAW**

The file system specified by *vfs\_id* is thawed. Modifications to the file system are still allowed after it is thawed, and the file system image might no longer be consistent after the thaw occurs. If the file system is not frozen at the time of the call, an error is returned. The *Argument* and *ArgumentSize* must both be 0.

The Journaled File System (JFS) supports only internal **fsctl** interfaces. Application programs should not call this function on a JFS file system, because **fsctl** is reserved for system management commands, such as the **chfs** command.

## **Parameters**

<i>vfs_id</i>	Identifies the file system to be acted upon. This information is returned by the <b>stat</b> subroutine in the <i>st_vfs</i> field of the <b>stat.h</b> file.
<i>Command</i>	Identifies the operation to be performed.
<i>Argument</i>	Specifies a pointer to a block of file system specific information that defines how the operation is to be performed.
<i>ArgumentSize</i>	Defines the size of the buffer pointed to by the <i>Argument</i> parameter.

## **Return Values**

Upon successful completion, the **fsctl** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## **Error Codes**

The **fsctl** subroutine fails if any of the following errors are true:

<b>EINVAL</b>	The <i>vfs_id</i> parameter does not identify a valid file system.
<b>EINVAL</b>	The <i>Command</i> parameter is not recognized by the file system.
<b>EINVAL</b>	The timeout specified to <b>FSCNTL_FREEZE</b> or <b>FSCNTL_REFREEZE</b> is invalid.
<b>EALREADY</b>	The <i>Command</i> parameter was <b>FSCNTL_FREEZE</b> and the file system specified was already frozen.
<b>EALREADY</b>	The <i>Command</i> parameter was <b>FSCNTL_REFREEZE</b> or <b>FSCNTL_THAW</b> and the file system specified was not frozen.

## **Related Information**

The **chfs** command.

The **stat.h** file.

## **fseek, fseeko, fseeko64, rewind, ftell, ftello, ftello64, fgetpos, fgetpos64, fsetpos, or fsetpos64 Subroutine**

### **Purpose**

Repositions the file pointer of a stream.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

```
#include <stdio.h>
```

```
int fseek ( Stream, Offset, Whence)
```

```
FILE *Stream;
```

```
long int Offset;
```

```
int Whence;
```

```
void rewind (Stream)
```

```
FILE *Stream;
```

```
long int ftell (Stream)
```

```
FILE *Stream;
```

```
int fgetpos (Stream, Position)
```

```
FILE *Stream;
```

```
fpos_t *Position;
```

```
int fsetpos (Stream, Position)
```

```
FILE *Stream;
```

```
const fpos_t *Position;
```

```
int fseeko ( Stream, Offset, Whence)
```

```
FILE *Stream;
```

```
off_t Offset;
```

```
int Whence;
```

```
int fseeko64 ( Stream, Offset, Whence)
```

```
FILE *Stream;
```

```
off64_t Offset;
```

```
int Whence;
```

```
off_t int ftello (Stream)
```

```
FILE *Stream;
```

```
off64_t int ftello64 (Stream)
```

```
FILE *Stream;
```

```
int fgetpos64 (Stream, Position)
```

```
FILE *Stream;
```

```
fpos64_t *Position;
```

```
int fsetpos64 (Stream, Position)
```

```
FILE *Stream;
```

```
const fpos64_t *Position;
```

## Description

The **fseek**, **fseeko** and **fseeko64** subroutines set the position of the next input or output operation on the I/O stream specified by the *Stream* parameter. The position if the next operation is determined by the *Offset* parameter, which can be either positive or negative.

The **fseek**, **fseeko** and **fseeko64** subroutines set the file pointer associated with the specified *Stream* as follows:

- If the *Whence* parameter is set to the **SEEK\_SET** value, the pointer is set to the value of the *Offset* parameter.
- If the *Whence* parameter is set to the **SEEK\_CUR** value, the pointer is set to its current location plus the value of the *Offset* parameter.
- If the *Whence* parameter is set to the **SEEK\_END** value, the pointer is set to the size of the file plus the value of the *Offset* parameter.

The **fseek**, **fseeko**, and **fseeko64** subroutine are unsuccessful if attempted on a file that has not been opened using the **fopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine. In particular, the **fseek** subroutine cannot be used on a terminal or on a file opened with the **popen** (“popen Subroutine” on page 1144) subroutine. The **fseek** and **fseeko** subroutines will also fail when the resulting offset is larger than can be properly returned.

The **rewind** subroutine is equivalent to calling the **fseek** subroutine using parameter values of (*Stream*,**SEEK\_SET**,**SEEK\_SET**), except that the **rewind** subroutine does not return a value.

The **fseek**, **fseeko**, **fseeko64** and **rewind** subroutines undo any effects of the **ungetc** and **ungetwc** subroutines and clear the end-of-file (EOF) indicator on the same stream.

The **fseek**, **fseeko**, and **fseeko64** function allows the file-position indicator to be set beyond the end of existing data in the file. If data is written later at this point, subsequent reads of data in the gap will return bytes of the value 0 until data is actually written into the gap.

A successful calls to the **fsetpos** or **fsetpos64** subroutines clear the **EOF** indicator and undoes any effects of the **ungetc** and **ungetwc** subroutines.

After an **fseek**, **fseeko**, **fseeko64** or a **rewind** subroutine, the next operation on a file opened for update can be either input or output.

**ftell**, **ftello** and **ftello64** subroutines return the position current value of the file-position indicator for the stream pointed to by the *Stream* parameter. **ftell** and **ftello** will fail if the resulting offset is larger than can be properly returned.

The **fgetpos** and **fgetpos64** subroutines store the current value of the file-position indicator for the stream pointed to by the *Stream* parameter in the object pointed to by the *Position* parameter. The **fsetpos** and **fsetpos64** set the file-position indicator for *Stream* according to the value of the *Position* parameter, which must be the result of a prior call to **fgetpos** or **fgetpos64** subroutine. **fgetpos** and **fsetpos** will fail if the resulting offset is larger than can be properly returned.

## Parameters

<i>Stream</i>	Specifies the input/output (I/O) stream.
<i>Offset</i>	Determines the position of the next operation.
<i>Whence</i>	Determines the value for the file pointer associated with the <i>Stream</i> parameter.
<i>Position</i>	Specifies the value of the file-position indicator.

## Return Values

Upon successful completion, the **fseek**, **fseeko** and **fseeko64** subroutine return a value of 0. Otherwise, it returns a value of -1.

Upon successful completion, the **ftell**, **ftello** and **ftello64** subroutine return the offset of the current byte relative to the beginning of the file associated with the named stream. Otherwise, a **long int** value of -1 is returned and the **errno** global variable is set.

Upon successful completion, the **fgetpos**, **fgetpos64**, **fsetpos** and **fsetpos64** subroutines return a value of 0. Otherwise, a nonzero value is returned and the **errno** global variable is set to the specific error.

The **errno** global variable is used to determine if an error occurred during a **rewind** subroutine call.

## Error Codes

If the **fseek**, **fseeko**, **fseeko64**, **ftell**, **ftello**, **ftello64** or **rewind** subroutine are unsuccessful because the stream is unbuffered or the stream buffer needs to be flushed and the call to the subroutine causes an underlying **lseek** or **write** subroutine to be invoked, it returns one or more of the following error codes:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> flag is set for the file descriptor, delaying the process in the write operation.
<b>EBADF</b>	Indicates that the file descriptor underlying the <i>Stream</i> parameter is not open for writing.
<b>EFBIG</b>	Indicates that an attempt has been made to write to a file that exceeds the file-size limit of the process or the maximum file size.
<b>EFBIG</b>	Indicates that the file is a regular file and that an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.
<b>EINTR</b>	Indicates that the write operation has been terminated because the process has received a signal, and either no data was transferred, or the implementation does not report partial transfers for this file.
<b>EIO</b>	Indicates that the process is a member of a background process group attempting to perform a <b>write</b> subroutine to its controlling terminal, the <b>TOSTOP</b> flag is set, the process is not ignoring or blocking the <b>SIGTTOU</b> signal, and the process group of the process is orphaned. This error may also be returned under implementation-dependent conditions.
<b>ENOSPC</b>	Indicates that no remaining free space exists on the device containing the file.
<b>EPIPE</b>	Indicates that an attempt has been made to write to a pipe or FIFO that is not open for reading by any process. A <b>SIGPIPE</b> signal will also be sent to the process.
<b>EINVAL</b>	Indicates that the <i>Whence</i> parameter is not valid. The resulting file-position indicator will be set to a negative value. The <b>EINVAL</b> error code does not apply to the <b>ftell</b> and <b>rewind</b> subroutines.
<b>ESPIPE</b>	Indicates that the file descriptor underlying the <i>Stream</i> parameter is associated with a pipe, FIFO, or socket.
<b>EOVERFLOW</b>	Indicates that for <i>fseek</i> , the resulting file offset would be a value that cannot be represented correctly in an object of type <i>long</i> .
<b>EOVERFLOW</b>	Indicates that for <i>fseeko</i> , the resulting file offset would be a value that cannot be represented correctly in an object of type <i>off_t</i> .
<b>ENXIO</b>	Indicates that a request was made of a non-existent device, or the request was outside the capabilities of the device.

The **fgetpos** and **fsetpos** subroutines are unsuccessful due to the following conditions:

<b>EINVAL</b>	Indicates that either the <i>Stream</i> or the <i>Position</i> parameter is not valid. The <b>EINVAL</b> error code does not apply to the <b>fgetpos</b> subroutine.
<b>EBADF</b>	Indicates that the file descriptor underlying the <i>Stream</i> parameter is not open for writing.
<b>ESPIPE</b>	Indicates that the file descriptor underlying the <i>Stream</i> parameter is associated with a pipe, FIFO, or socket.



The **fseek**, **fseeko**, **ftell**, **ftello**, **fgetpos**, and **fsetpos** subroutines are unsuccessful under the following condition:

**EOVERFLOW**

The resulting could not be returned properly.

## Related Information

The **closedir** (“**opendir**, **readdir**, **tell**, **seekdir**, **rewinddir**, **closedir**, **opendir64**, **readdir64**, **tell64**, **seek64**, **rewind64**, or **closedir64** Subroutine” on page 943) subroutine, **fopen**, **fopen64**, **freopen**, **freopen64** or **fdopen** (“**fopen**, **fopen64**, **freopen**, **freopen64** or **fdopen** Subroutine” on page 288) subroutine, **lseek** or **lseek64** (“**lseek**, **llseek** or **llseek64** Subroutine” on page 765) subroutine, **opendir**, **readdir**, **rewinddir**, **seekdir**, or **tell** (“**opendir**, **readdir**, **tell**, **seekdir**, **rewinddir**, **closedir**, **opendir64**, **readdir64**, **tell64**, **seek64**, **rewind64**, or **closedir64** Subroutine” on page 943) subroutine, **popen** (“**popen** Subroutine” on page 1144) subroutine, **ungetc** or **ungetwc** subroutine, **write**, **writex**, **writev**, or **writevx** subroutine.

Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fsync or fsync\_range Subroutine

### Purpose

Writes changes in a file to permanent storage.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int fsync ( FileDescriptor )
int FileDescriptor;
```

```
int fsync_range ( FileDescriptor, how, start, length )
int FileDescriptor;
int how;
off_t start;
off_t length;
```

### Description

The **fsync** subroutine causes all modified data in the open file specified by the *FileDescriptor* parameter to be saved to permanent storage. On return from the **fsync** subroutine, all updates have been saved on permanent storage.

The **fsync\_range** subroutine causes all modified data in the specified range of the open file specified by the *FileDescriptor* parameter to be saved to permanent storage. On return from the **fsync\_range** subroutine, all updates in the specified range have been saved on permanent storage.

Data written to a file that a process has opened for deferred update (with the **O\_DEFER** flag) is not written to permanent storage until another process issues an **fsync\_range** or **fsync** call against this file or runs a synchronous **write** subroutine (with the **O\_SYNC** flag) on this file. See the **fcntl.h** file and the **open** subroutine for descriptions of the **O\_DEFER** and **O\_SYNC** flags respectively.

**Note:** The file identified by the *FileDescriptor* parameter must be open for writing when the **fsync\_range** or **fsync** subroutine is issued or the call is unsuccessful. This restriction was not enforced in BSD systems.

## Parameters

<i>FileDescriptor</i>	A valid, open file descriptor.
<i>how</i>	How to flush, <b>O_DSYNC</b> , <b>O_NOCACHE</b> or <b>O_SYNC</b> .
	<b>O_DSYNC</b> Write file data and metadata to retrieve the data for the specified range.
	<b>O_NOCACHE</b> Write the data in the range and release full memory pages in the byte range. The data will no longer be cached.
	<b>O_SYNC</b> Write all modified file data and metadata for the specified range.
<i>start</i>	Starting file offset.
<i>length</i>	Length, or zero for everything.

## Return Values

Upon successful completion, the **fsync** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

Upon successful completion, the **fsync\_range** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **fsync** subroutine is unsuccessful if one or more of the following are true:

<b>EIO</b>	An I/O error occurred while reading from or writing to the file system.
<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor open for writing.
<b>EINVAL</b>	The file is not a regular file.
<b>EINTR</b>	The <b>fsync</b> subroutine was interrupted by a signal.

## Related Information

The **open**, **openx**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **sync** subroutine, **write**, **writex**, **writev**, or **writevx** subroutine.

The **fcntl.h** file.

Files, Directories, and File Systems Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* contains information about i-nodes, file descriptors, file-space allocation, and more.

---

## ftok Subroutine

### Purpose

Generates a standard interprocess communication key.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/types.h>
#include <sys/ipc.h>
```

```
key_t ftok ( Path, ID)
char *Path;
int ID;
```

## Description

**Attention:** If the *Path* parameter of the **ftok** subroutine names a file that has been removed while keys still refer to it, the **ftok** subroutine returns an error. If that file is then re-created, the **ftok** subroutine will probably return a key different from the original one.

**Attention:** Each installation should define standards for forming keys. If standards are not adhered to, unrelated processes may interfere with each other's operation.

**Attention:** The **ftok** subroutine does not guarantee unique key generation. However, the occurrence of key duplication is very rare and mostly for across file systems.

The **ftok** subroutine returns a key, based on the *Path* and *ID* parameters, to be used to obtain interprocess communication identifiers. The **ftok** subroutine returns the same key for linked files if called with the same *ID* parameter. Different keys are returned for the same file if different *ID* parameters are used.

All interprocess communication facilities require you to supply a key to the **msgget**, **semget**, and **shmget** subroutines in order to obtain interprocess communication identifiers. The **ftok** subroutine provides one method for creating keys, but other methods are possible. For example, you can use the project ID as the most significant byte of the key, and use the remaining portion as a sequence number.

## Parameters

*Path* Specifies the path name of an existing file that is accessible to the process.  
*ID* Specifies a character that uniquely identifies a project.

## Return Values

When successful, the **ftok** subroutine returns a key that can be passed to the **msgget**, **semget**, or **shmget** subroutine.

## Error Codes

The **ftok** subroutine returns the value **(key\_t)-1** if one or more of the following are true:

- The file named by the *Path* parameter does not exist.
- The file named by the *Path* parameter is not accessible to the process.
- The *ID* parameter has a value of 0.

## Related Information

The **msgget** (“msgget Subroutine” on page 881) subroutine, **semget** subroutine, **shmget** subroutine.

Subroutines Overview and Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ftw or ftw64 Subroutine

### Purpose

Walks a file tree.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <ftw.h>
```

```
int ftw ( Path, Function, Depth)
char *Path;
int (*Function)(const char*, const struct stat*, int);
int Depth;
```

```
int ftw64 ( Path, Function, Depth)
char *Path;
int (*Function)(const char*, const struct stat64*, int);
int Depth;
```

### Description

The **ftw** and **ftw64** subroutines recursively searches the directory hierarchy that descends from the directory specified by the *Path* parameter.

For each file in the hierarchy, the **ftw** and **ftw64** subroutines call the function specified by the *Function* parameter. **ftw** passes it a pointer to a null-terminated character string containing the name of the file, a pointer to a *stat* structure containing information about the file, and an integer. **ftw64** passes it a pointer to a null-terminated character string containing the name of the file, a pointer to a **stat64** structure containing information about the file, and an integer.

The integer passed to the *Function* parameter identifies the file type with one of the following values:

<b>FTW_F</b>	Regular file
<b>FTW_D</b>	Directory
<b>FTW_DNR</b>	Directory that cannot be read
<b>FTW_SL</b>	Symbolic Link
<b>FTW_NS</b>	File for which the <b>stat</b> structure could not be executed successfully

If the integer is **FTW-DNR**, the files and subdirectories contained in that directory are not processed.

If the integer is **FTW-NS**, the **stat** structure contents are meaningless. An example of a file that causes **FTW-NS** to be passed to the *Function* parameter is a file in a directory for which you have read permission but not execute (search) permission.

The **ftw** and **ftw64** subroutines finish processing a directory before processing any of its files or subdirectories.

The **ftw** and **ftw64** subroutines continue the search until the directory hierarchy specified by the *Path* parameter is completed, an invocation of the function specified by the *Function* parameter returns a nonzero value, or an error is detected within the **ftw** and **ftw64** subroutines, such as an I/O error.

The **ftw** and **ftw64** subroutines traverse symbolic links encountered in the resolution of the *Path* parameter, including the final component. Symbolic links encountered while walking the directory tree rooted at the *Path* parameter are not traversed.

The **ftw** and **ftw64** subroutines use one file descriptor for each level in the tree. The *Depth* parameter specifies the maximum number of file descriptors to be used. In general, the **ftw** and **ftw64** subroutines runs faster if the value of the *Depth* parameter is at least as large as the number of levels in the tree. However, the value of the *Depth* parameter must not be greater than the number of file descriptors currently available for use. If the value of the *Depth* parameter is 0 or a negative number, the effect is the same as if it were 1.

Because the **ftw** and **ftw64** subroutines are recursive, it is possible for it to terminate with a memory fault due to stack overflow when applied to very deep file structures.

The **ftw** and **ftw64** subroutines use the **malloc** subroutine to allocate dynamic storage during its operation. If the **ftw** and **ftw64** subroutines is terminated prior to its completion, such as by the **longjmp** subroutine being executed by the function specified by the *Function* parameter or by an interrupt routine, the **ftw** and **ftw64** subroutines cannot free that storage. The storage remains allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have the function specified by the *Function* parameter return a nonzero value the next time it is called.

## Parameters

<i>Path</i>	Specifies the directory hierarchy to be searched.
<i>Function</i>	Specifies the file type.
<i>Depth</i>	Specifies the maximum number of file descriptors to be used. <i>Depth</i> cannot be greater than OPEN_MAX which is described in the sys/limits.h header file.

## Return Values

If the tree is exhausted, the **ftw** and **ftw64** subroutines returns a value of 0. If the subroutine pointed to by **fn** returns a nonzero value, **ftw** and **ftw64** subroutines stops its tree traversal and returns whatever value was returned by the subroutine pointed to by **fn**. If the **ftw** and **ftw64** subroutines detects an error, it returns a -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **ftw** or **ftw64** subroutines detect an error, a value of -1 is returned and the **errno** global variable is set to indicate the error.

The **ftw** and **ftw64** subroutine are unsuccessful if:

<b>EACCES</b>	Search permission is denied for any component of the <i>Path</i> parameter or read permission is denied for <i>Path</i> .
<b>ENAMETOOLONG</b>	The length of the path exceeds <b>PATH_MAX</b> while <b>_POSIX_NO_TRUNC</b> is in effect.
<b>ENOENT</b>	The <i>Path</i> parameter points to the name of a file that does not exist or points to an empty string.
<b>ENOTDIR</b>	A component of the <i>Path</i> parameter is not a directory.

The **ftw** subroutine is unsuccessful if:

<b>EOVERFLOW</b>	A file in <i>Path</i> is of a size larger than 2 Gigabytes.
------------------	---

## Related Information

The **malloc**, **free**, **realloc**, **calloc**, **malloc**, **mallinfo**, or **alloca** (“malloc, free, realloc, calloc, malloc, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778) subroutine, **setjmp** or **longjmp** subroutine, **signal** subroutine, **stat** subroutine.

Searching and Sorting Example Program and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## fwide Subroutine

### Purpose

Set stream orientation.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <stdio.h>
#include <wchar.h>
int fwide (FILE * stream, int mode),
```

### Description

The **fwide** function determines the orientation of the stream pointed to by *stream*. If *mode* is greater than zero, the function first attempts to make the stream wide-oriented. If *mode* is less than zero, the function first attempts to make the stream byte-oriented. Otherwise, *mode* is zero and the function does not alter the orientation of the stream.

If the orientation of the stream has already been determined, **fwide** does not change it.

Because no return value is reserved to indicate an error, an application wishing to check for error situations should set *errno* to 0, then call **fwide**, then check *errno* and if it is non-zero, assume an error has occurred.

A call to **fwide** with *mode* set to zero can be used to determine the current orientation of a stream.

### Return Values

The **fwide** function returns a value greater than zero if, after the call, the stream has wide-orientation, a value less than zero if the stream has byte-orientation, or zero if the stream has no orientation.

### Errors

The **fwide** function may fail if:

**EBADF**                    The stream argument is not a valid stream.

## Related Information

The **wchar.h** file

---

## fwprintf, wprintf, swprintf Subroutines

### Purpose

Print formatted wide-character output.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <stdio.h>
#include <wchar.h>

int fwprintf ( FILE * stream, const wchar_t * format, . . . )
int wprintf (const wchar_t * format, . . . )
int swprintf (wchar_t *s, size_t n, const wchar_t * format, . . . )
```

### Description

The **fwprintf** function places output on the named output stream. The **wprintf** function places output on the standard output stream **stdout**. The **swprintf** function places output followed by the null wide-character in consecutive wide-characters starting at **\*s**; no more than **n** wide-characters are written, including a terminating null wide-character, which is always added (unless **n** is zero).

Each of these functions converts, formats and prints its arguments under control of the **format** wide-character string. The **format** is composed of zero or more directives: **ordinary wide-characters**, which are simply copied to the output stream and **conversion specifications**, each of which results in the fetching of zero or more arguments. The results are undefined if there are insufficient arguments for the **format**. If the **format** is exhausted while arguments remain, the excess arguments are evaluated but are otherwise ignored.

EX Conversions can be applied to the **nth** argument after the **format** in the argument list, rather than to the next unused argument. In this case, the conversion wide-character % (see below) is replaced by the sequence **%n\$**, where **n** is a decimal integer in the range [1, {NL\_ARGMAX}], giving the position of the argument in the argument list. This feature provides for the definition of format wide-character strings that select arguments in an order appropriate to specific languages (see the EXAMPLES section).

In format wide-character strings containing the **%n\$** form of conversion specifications, numbered arguments in the argument list can be referenced from the format wide-character string as many times as required.

In format wide-character strings containing the % form of conversion specifications, each argument in the argument list is used exactly once.

All forms of the **fwprintf** functions allow for the insertion of a language-dependent radix character in the output string, output as a wide-character value. The radix character is defined in the program's locale (category LC\_NUMERIC). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (.).

EX Each conversion specification is introduced by the % wide-character or by the wide-character sequence **%n\$**, after which the following appear in sequence:

- Zero or more **flags** (in any order), which modify the meaning of the conversion specification.
- An optional minimum **field width**. If the converted value has fewer wide-characters than the field width, it will be padded with spaces by default on the left; it will be padded on the right, if the left-adjustment flag (-), described below, is given to the field width. The field width takes the form of an asterisk (\*), described below, or a decimal integer.

- An optional **precision** that gives the minimum number of digits to appear for the **d**, **i**, **o**, **u**, **x** and **X** conversions; the number of digits to appear after the radix character for the **e**, **E** and **f** conversions; the maximum number of significant digits for the **g** and **G** conversions; or the maximum number of wide-characters to be printed from a string in **s** conversions. The precision takes the form of a period (.) followed either by an asterisk (\*), described below, or an optional decimal digit string, where a null digit string is treated as 0. If a precision appears with any other conversion wide-character, the behaviour is undefined.
- An optional **I** (lowercase **L**), **L**, **h**, **H**, **D** or **DD** specifies one of the following conversions:
  - An optional **I** specifying that a following **c** conversion wide-character applies to a **wint\_t** argument.
  - An optional **I** specifying that a following **s** conversion wide-character applies to a **wchar\_t** argument.
  - An optional **I** specifying that a following **d**, **i**, **o**, **u**, **x** or **X** conversion wide-character applies to a type **long int** or **unsigned long int** argument.
  - An optional **I** specifying that a following **n** conversion wide-character applies to a pointer to a type **long int** argument.
  - An optional **L** specifying that a following **e**, **E**, **f**, **g** or **G** conversion wide-character applies to a type **long double** argument.
  - An optional **h** specifying that a following **d**, **i**, **o**, **u**, **x** or **X** conversion wide-character applies to a type **short int** or type **unsigned short int** argument (the argument that will be promoted according to the integral promotions, and its value will be converted to type **short int** or **unsigned short int** before printing).
  - An optional **h** specifying that a following **n** conversion wide-character applies to a pointer to a type **short int** argument.
  - An optional **H** specifying that a following **e**, **E**, **f**, **g**, or **G** conversion wide-character applies to a **\_Decimal32** parameter.
  - An optional **D** specifying that a following **e**, **E**, **f**, **g**, or **G** conversion wide-character applies to a **\_Decimal64** parameter.
  - An optional **DD** specifying that a following **e**, **E**, **f**, **g**, or **G** conversion wide-character applies to a **\_Decimal128** parameter.

If an **I**, **L**, **h**, **H**, **D**, or **DD** appears with any other conversion wide-character, the behavior is undefined.

- A **conversion wide-character** that indicates the type of conversion to be applied.

A field width, or precision, or both, may be indicated by an asterisk (\*). In this case an argument of type **int** supplies the field width or precision. Arguments specifying field width, or precision, or both must appear in that order before the argument, if any, to be converted. A negative field width is taken as a - flag followed by a positive field width. A negative precision is taken as if EX the precision were omitted. In format wide-character strings containing the **%n\$** form of a conversion specification, a field width or precision may be indicated by the sequence **\*m\$**, where **m** is a decimal integer in the range [1, {NL\_ARGMAX}] giving the position in the argument list (after the format argument) of an integer argument containing the field width or precision, for example:

```
wprintf(L"%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```

The **format** can contain either numbered argument specifications (that is, **%n\$** and **\*m\$**), or unnumbered argument specifications (that is, **%** and **\***), but normally not both. The only exception to this is that **%%** can be mixed with the **%n\$** form. The results of mixing numbered and unnumbered argument specifications in a **format** wide-character string are undefined. When numbered argument specifications are used, specifying the **N**th argument requires that all the leading arguments, from the first to the (N-1)th, are specified in the format wide-character string.



The flag wide-characters and their meanings are:

'	The integer portion of the result of a decimal conversion (%i, %d, %u, %f, %g or %G) will be formatted with thousands' grouping wide-characters. For other conversions the behaviour is undefined. The non-monetary grouping wide-character is used.
-	The result of the conversion will be left-justified within the field. The conversion will be right-justified if this flag is not specified.
+	The result of a signed conversion will always begin with a sign (+ or -). The conversion will begin with a sign only when a negative value is converted if this flag is not specified.
space	If the first wide-character of a signed conversion is not a sign or if a signed conversion results in no wide-characters, a space will be prefixed to the result. This means that if the space and + flags both appear, the space flag will be ignored.
#	This flag specifies that the value is to be converted to an alternative form. For o conversion, it increases the precision (if necessary) to force the first digit of the result to be 0. For x or X conversions, a non-zero result will have 0x (or 0X) prefixed to it. For e, E, f, g or G conversions, the result will always contain a radix character, even if no digits follow it. Without this flag, a radix character appears in the result of these conversions only if a digit follows it. For g and G conversions, trailing zeros will <b>not</b> be removed from the result as they normally are. For other conversions, the behavior is undefined.
0	For d, i, o, u, x, X, e, E, f, g and G conversions, leading zeros (following any indication of sign or base) are used to pad to the field width; no space padding is performed. If the 0 and - flags both appear, the 0 flag will be ignored. For d, i, o, u, x and X conversions, if a precision is specified, the 0 flag will be ignored. If the 0 and ' flags both appear, the grouping wide-characters are inserted before zero padding. For other conversions, the behavior is undefined.

The conversion wide-characters and their meanings are:

d,i	The <b>int</b> argument is converted to a signed decimal in the style [-] <b>dddd</b> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.
o	The <b>unsigned int</b> argument is converted to unsigned octal format in the style <b>dddd</b> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.
u	The <b>unsigned int</b> argument is converted to unsigned decimal format in the style <b>dddd</b> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.
x	The <b>unsigned int</b> argument is converted to unsigned hexadecimal format in the style <b>dddd</b> ; the letters abcdef are used. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.
X	Behaves the same as the x conversion wide-character except that letters ABCDEF are used instead of abcdef.
f	The <b>double</b> argument is converted to decimal notation in the style [-] <b>ddd.ddd</b> , where the number of digits after the radix character is equal to the precision specification. If the precision is missing, it is taken as 6; if the precision is explicitly 0 and no # flag is present, no radix character appears. If a radix character appears, at least one digit appears before it. The value is rounded to the appropriate number of digits. The <b>fwprintf</b> family of functions may make available wide-character string representations for infinity and NaN.
e, E	The <b>double</b> argument is converted in the style [-] <b>d.ddde +/- dd</b> , where there is one digit before the radix character (which is non-zero if the argument is non-zero) and the number of digits after it is equal to the precision; if the precision is missing, it is taken as 6; if the precision is 0 and no # flag is present, no radix character appears. The value is rounded to the appropriate number of digits. The E conversion wide-character will produce a number with E instead of e introducing the exponent. The exponent always contains at least two digits. If the value is 0, the exponent is 0.

	The <b>fwprintf</b> family of functions may make available wide-character string representations for infinity and NaN.
<b>g, G</b>	The <b>double</b> argument is converted in the style <b>f</b> or <b>e</b> (or in the style <b>E</b> in the case of a <b>G</b> conversion wide-character), with the precision specifying the number of significant digits. If an explicit precision is 0, it is taken as 1. The style used depends on the value converted; style <b>e</b> (or <b>E</b> ) will be used only if the exponent resulting from such a conversion is less than -4 or greater than or equal to the precision. Trailing zeros are removed from the fractional portion of the result; a radix character appears only if it is followed by a digit. The <b>fwprintf</b> family of functions may make available wide-character string representations for infinity and NaN.
<b>c</b>	If no <b>l</b> (ell) qualifier is present, the <b>int</b> argument is converted to a wide-character as if by calling the <b>btowc</b> function and the resulting wide-character is written. Otherwise the <b>wint_t</b> argument is converted to <b>wchar_t</b> , and written.
<b>s</b>	If no <b>l</b> (ell) qualifier is present, the argument must be a pointer to a character array containing a character sequence beginning in the initial shift state. Characters from the array are converted as if by repeated calls to the <b>mbrtowc</b> function, with the conversion state described by an <b>mbstate_t</b> object initialised to zero before the first character is converted, and written up to (but not including) the terminating null wide-character. If the precision is specified, no more than that many wide-characters are written. If the precision is not specified or is greater than the size of the array, the array must contain a null wide-character. If an <b>l</b> (ell) qualifier is present, the argument must be a pointer to an array of type <b>wchar_t</b> . Wide characters from the array are written up to (but not including) a terminating null wide-character. If no precision is specified or is greater than the size of the array, the array must contain a null wide-character. If a precision is specified, no more than that many wide-characters are written.
<b>p</b>	The argument must be a pointer to void. The value of the pointer is converted to a sequence of printable wide-characters, in an implementation-dependent manner. The argument must be a pointer to an integer into which is written the number of wide-characters written to the output so far by this call to one of the <b>fwprintf</b> functions. No argument is converted.
<b>C</b>	Same as <b>lc</b> .
<b>S</b>	Same as <b>ls</b> .
<b>%</b>	Output a <b>%</b> wide-character; no argument is converted. The entire conversion specification must be <b>%%</b> .

If a conversion specification does not match one of the above forms, the behavior is undefined.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Characters generated by **fwprintf** and **wprintf** are printed as if **fputwc** had been called.

The **st\_ctime** and **st\_mtime** fields of the file will be marked for update between the call to a successful execution of **fwprintf** or **wprintf** and the next successful completion of a call to **fflush** or **fclose** on the same stream or a call to **exit** or **abort**.

## Return Values

Upon successful completion, these functions return the number of wide-characters transmitted excluding the terminating null wide-character in the case of **swprintf** or a negative value if an output error was encountered.

## Error Codes

For the conditions under which **fwprintf** and **wprintf** will fail and may fail, refer to **fputwc**. In addition, all forms of **fwprintf** may fail if:

<b>EILSEQ</b>	A wide-character code that does not correspond to a valid character has been detected
<b>EINVAL</b>	There are insufficient arguments. In addition, <b>wprintf</b> and <b>fwprintf</b> may fail if:
<b>ENOMEM</b>	Insufficient storage space is available.

## Examples

To print the language-independent date and time format, the following statement could be used:

```
wprintf (format, weekday, month, day, hour, min);
```

For American usage, format could be a pointer to the wide-character string:

```
L"%s, %s %d, %d:%.2d\n"
```

producing the message:

```
Sunday, July 3, 10:02
```

whereas for German usage, format could be a pointer to the wide-character string:

```
L"%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

producing the message:

```
Sonntag, 3. July, 10:02
```

## Related Information

The **btowc** (“btowc Subroutine” on page 124) subroutine, **fputwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **fwscanf** (“fwscanf, wscanf, swscanf Subroutines”) subroutine, **setlocale** subroutine, **mbrtowc** (“mbrtowc Subroutine” on page 793) subroutine.

The **wchar.h** file.

---

## fwscanf, wscanf, swscanf Subroutines

### Purpose

Convert formatted wide-character input.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <stdio.h>
#include <wchar.h>

int fwscanf (FILE * stream, const wchar_t * format, ...);
int wscanf (const wchar_t * format, ...);
int swscanf (const wchar_t * s, const wchar_t * format, ...);
```

### Description

The **fwscanf** function reads from the named input stream. The **wscanf** function reads from the standard input stream `stdin`. The **swscanf** function reads from the wide-character string `s`. Each function reads wide-characters, interprets them according to a format, and stores the results in its arguments. Each expects, as arguments, a control wide-character string format described below, and a set of pointer arguments indicating where the converted input should be stored. The result is undefined if there are insufficient arguments for the format. If the format is exhausted while arguments remain, the excess arguments are evaluated but are otherwise ignored.

Conversions can be applied to the **n**th argument after the **format** in the argument list, rather than to the next unused argument. In this case, the conversion wide-character `%` (see below) is replaced by the sequence `%n$`, where **n** is a decimal integer in the range `[1, {NL_ARGMAX}]`. This feature provides for the

definition of format wide-character strings that select arguments in an order appropriate to specific languages. In format wide-character strings containing the `%n$` form of conversion specifications, it is unspecified whether numbered arguments in the argument list can be referenced from the format wide-character string more than once.

The format can contain either form of a conversion specification, that is, `%` or `%n$`, but the two forms cannot normally be mixed within a single format wide-character string. The only exception to this is that `%%` or `%*` can be mixed with the `%n$` form.

The **fwscanf** function in all its forms allows for detection of a language-dependent radix character in the input string, encoded as a wide-character value. The radix character is defined in the program's locale (category `LC_NUMERIC`). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (`.`).

The format is a wide-character string composed of zero or more directives. Each directive is composed of one of the following: one or more white-space wide-characters (space, tab, newline, vertical-tab or form-feed characters); an ordinary wide-character (neither `%` nor a white-space character); or a conversion specification. Each conversion specification is introduced by a `%` or the sequence `%n$` after which the following appear in sequence:

- An optional assignment-suppressing character `*`.
- An optional non-zero decimal integer that specifies the maximum field width.
- An optional size modifier **h**, **H**, **I** (lowercase *L*), **L**, **D**, or **DD** indicating the size of the receiving object.
  - Must precede the **c**, **s** and **[** conversion wide-characters with the **I** (lowercase *L*) if the corresponding argument is a pointer to **wchar\_t**.
  - Must precede the **d**, **i** and **n** conversion wide-characters with the **h** if the corresponding argument is a pointer to **short int** or with the **I** (lowercase *L*) if it is a pointer to **long int**.
  - Must precede the **o**, **u** and **x** conversion wide-characters with the **h** if the corresponding argument is a pointer to **unsigned short int** or with **I** (lowercase *L*) if it is a pointer to **unsigned long int**.
  - Must precede the **e**, **f** and **g** conversion wide-characters with **I** (lowercase *L*) if the corresponding argument is a pointer to **double** or with the **L** if it is a pointer to long double.
  - Must precede the **e**, **f** and **g** conversion wide-characters with the **H** if the corresponding argument is a pointer to **\_Decimal32**.
  - Must precede the **e**, **f** and **g** conversion wide-characters with the **D** if the corresponding argument is a pointer to **\_Decimal64**.
  - Must precede the **e**, **f** and **g** conversion wide-characters with the **DD** if the corresponding argument is a pointer to **\_Decimal128**.

If an **I** (lowercase *L*), **L**, **h**, **H**, **D**, or **DD** appears with any other *L* conversion wide-character, the behavior is undefined.

- A conversion wide-character that specifies the type of conversion to be applied. The valid conversion wide-characters are described below.

The **fwscanf** functions execute each directive of the format in turn. If a directive fails, as detailed below, the function returns. Failures are described as input failures (due to the unavailability of input bytes) or matching failures (due to inappropriate input).

A directive composed of one or more white-space wide-characters is executed by reading input until no more valid input can be read, or up to the first wide-character which is not a white-space wide-character, which remains unread.

A directive that is an ordinary wide-character is executed as follows. The next wide-character is read from the input and compared with the wide-character that comprises the directive; if the comparison shows that they are not equivalent, the directive fails, and the differing and subsequent wide-characters remain unread.

A directive that is a conversion specification defines a set of matching input sequences, as described below for each conversion wide-character. A conversion specification is executed in the following steps:

Input white-space wide-characters (as specified by **iswspace**) are skipped, unless the conversion specification includes a `]`, `c` or `n` conversion character.

An item is read from the input, unless the conversion specification includes an `n` conversion wide-character. An input item is defined as the longest sequence of input wide-characters, not exceeding any specified field width, which is an initial subsequence of a matching sequence. The first wide-character, if any, after the input item remains unread. If the length of the input item is 0, the execution of the conversion specification fails; this condition is a matching failure, unless end-of-file, an encoding error, or a read error prevented input from the stream, in which case it is an input failure.

Except in the case of a `%` conversion wide-character, the input item (or, in the case of a `%n` conversion specification, the count of input wide-characters) is converted to a type appropriate to the conversion wide-character. If the input item is not a matching sequence, the execution of the conversion specification fails; this condition is a matching failure. Unless assignment suppression was indicated by a `*`, the result of the conversion is placed in the object pointed to by the first argument following the **format** argument that has not already received a conversion result if the conversion specification is introduced by `%`, or in the `n`th argument if introduced by the wide-character sequence `%n$`. If this object does not have an appropriate type, or if the result of the conversion cannot be represented in the space provided, the behavior is undefined.

The following conversion wide-characters are valid:

- d** Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of **wcstol** with the value 10 for the base argument. In the absence of a size modifier, the corresponding argument must be a pointer to **int**.
- i** Matches an optionally signed integer, whose format is the same as expected for the subject sequence of **wcstol** with 0 for the base argument. In the absence of a size modifier, the corresponding argument must be a pointer to **int**.
- o** Matches an optionally signed octal integer, whose format is the same as expected for the subject sequence of **wcstoul** with the value 8 for the base argument. In the absence of a size modifier, the corresponding argument must be a pointer to **unsigned int**.
- u** Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of **wcstoul** with the value 10 for the base argument. In the absence of a size modifier, the corresponding argument must be a pointer to **unsigned int**.
- x** Matches an optionally signed hexadecimal integer, whose format is the same as expected for the subject sequence of **wcstoul** with the value 16 for the base argument. In the absence of a size modifier, the corresponding argument must be a pointer to **unsigned int**.
- e, f, g** Matches an optionally signed floating-point number, whose format is the same as expected for the subject sequence of **wcstod**. In the absence of a size modifier, the corresponding argument must be a pointer to **float**.  
If the **fwprintf** family of functions generates character string representations for infinity and NaN (a 7858 symbolic entity encoded in floating-point format) to support the ANSI/IEEE Std 754:1985 standard, the **fwscanf5** family of functions will recognise them as input.
- s** Matches a sequence of non white-space wide-characters. If no `l` (`ell`) qualifier is present, characters from the input field are converted as if by repeated calls to the **wcrtomb** function, with the conversion state described by an **mbstate\_t** object initialised to zero before the first wide-character is converted. The corresponding argument must be a pointer to a character array large enough to accept the sequence and the terminating null character, which will be added automatically.  
Otherwise, the corresponding argument must be a pointer to an array of **wchar\_t** large enough to accept the sequence and the terminating null wide-character, which will be added automatically.

- [ Matches a non-empty sequence of wide-characters from a set of expected wide-characters (the **scanset**). If no l (ell) qualifier is present, wide-characters from the input field are converted as if by repeated calls to the **wcrtomb** function, with the conversion state described by an **mbstate\_t** object initialised to zero before the first wide-character is converted. The corresponding argument must be a pointer to a character array large enough to accept the sequence and the terminating null character, which will be added automatically.  
If an l (ell) qualifier is present, the corresponding argument must be a pointer to an array of **wchar\_t** large enough to accept the sequence and the terminating null wide-character, which will be added automatically.  
The conversion specification includes all subsequent wide characters in the **format** string up to and including the matching right square bracket (]). The wide-characters between the square brackets (the **scanlist**) comprise the scanset, unless the wide-character after the left square bracket is a circumflex (^), in which case the scanset contains all wide-characters that do not appear in the scanlist between the circumflex and the right square bracket. If the conversion specification begins with [ ] or [^], the right square bracket is included in the scanlist and the next right square bracket is the matching right square bracket that ends the conversion specification; otherwise the first right square bracket is the one that ends the conversion specification. If a - is in the scanlist and is not the first wide-character, nor the second where the first wide-character is a ^;, nor the last wide-character, the behavior is implementation-dependent.
- c Matches a sequence of wide-characters of the number specified by the field width (1 if no field width is present in the conversion specification). If no l (ell) qualifier is present, wide-characters from the input field are converted as if by repeated calls to the **wcrtomb** function, with the conversion state described by an **mbstate\_t** object initialised to zero before the first wide-character is converted. The corresponding argument must be a pointer to a character array large enough to accept the sequence. No null character is added.  
Otherwise, the corresponding argument must be a pointer to an array of **wchar\_t** large enough to accept the sequence. No null wide-character is added.
- p Matches an implementation-dependent set of sequences, which must be the same as the set of sequences that is produced by the %p conversion of the corresponding **fwprintf** functions. The corresponding argument must be a pointer to a pointer to void. The interpretation of the input item is implementation-dependent. If the input item is a value converted earlier during the same program execution, the pointer that results will compare equal to that value; otherwise the behavior of the %p conversion is undefined.
- n No input is consumed. The corresponding argument must be a pointer to the integer into which is to be written the number of wide-characters read from the input so far by this call to the **fwscanf** functions. Execution of a %n conversion specification does not increment the assignment count returned at the completion of execution of the function.
- C Same as lc.
- S Same as ls.
- % Matches a single %; no conversion or assignment occurs. The complete conversion specification must be %%.

If a conversion specification is invalid, the behavior is undefined.

The conversion characters E, G and X are also valid and behave the same as, respectively, e, g and x.

If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before any wide-characters matching the current conversion specification (except for %n) have been read (other than leading white-space, where permitted), execution of the current conversion specification terminates with an input failure. Otherwise, unless execution of the current conversion specification is terminated with a matching failure, execution of the following conversion specification (if any) is terminated with an input failure.

Reaching the end of the string in **swscanf** is equivalent to encountering end-of-file for **fwscanf**.

If conversion terminates on a conflicting input, the offending input is left unread in the input. Any trailing white space (including newline) is left unread unless matched by a conversion specification. The success of literal matches and suppressed assignments is only directly determinable via the %n conversion specification.

The **fwscanf** and **wscanf** functions may mark the **st\_atime** field of the file associated with stream for update. The **st\_atime** field will be marked for update by the first successful execution of **fgetc**, **fgetwc**, **fgets**, **fgetws**, **fread**, **getc**, **getwc**, **getchar**, **getwchar**, **gets**, **fscanf** or **fwscanf** using stream that returns data not supplied by a prior call to **ungetc**.

In format strings containing the % form of conversion specifications, each argument in the argument list is used exactly once.

## Return Values

Upon successful completion, these functions return the number of successfully matched and assigned input items; this number can be 0 in the event of an early matching failure. If the input ends before the first matching failure or conversion, EOF is returned. If a read error occurs the error indicator for the stream is set, EOF is returned, and errno is set to indicate the error.

## Error Codes

For the conditions under which the **fwscanf** functions will fail and may fail, refer to **fgetwc**. In addition, **fwscanf** may fail if:

<b>EILSEQ</b>	Input byte sequence does not form a valid character.
<b>EINVAL</b>	There are insufficient arguments.

## Examples

The call:

```
int i, n; float x; char name[50];
n = wscanf(L"%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 Hamster
```

will assign to n the value 3, to i the value 25, to x the value 5.432, and name will contain the string Hamster.

The call:

```
int i; float x; char name[50];
(void) wscanf(L"%2d%f*d %[0123456789]", &i, &x, name);
```

with input:

```
56789 0123 56a72
```

will assign 56 to i, 789.0 to x, skip 0123, and place the string 56\0 in **name**. The next call to **getchar** will return the character a.

## Related Information

The **getwc** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **fwprintf** (“fwprintf, wprintf, swprintf Subroutines” on page 327) subroutine, **setlocale** subroutine, **wcstod** subroutine, **wcstol** subroutine, **wcstoul** subroutine, **wctomb** subroutine.

The **wchar.h** file.

---

## gai\_strerror Subroutine

### Purpose

Facilitates consistent error information from EAI\_\* values returned by the getaddrinfo subroutine.

### Library

Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <netdb.h>
char *
gai_strerror (ecode)
int ecode;
int
gai_strerror_r (ecode, buf, buflen)
int ecode;
char *buf;
int buflen;
```

### Description

For multithreaded environments, the second version should be used. In **gai\_strerror\_r**, *buf* is a pointer to a data area to be filled in. *buflen* is the length (in bytes) available in *buf*.

It is the caller's responsibility to insure that *buf* is sufficiently large to store the requested information, including a trailing null character. It is the responsibility of the function to insure that no more than *buflen* bytes are written into *buf*.

### Return Values

If successful, a pointer to a string containing an error message appropriate for the EAI\_\* errors is returned. If *ecode* is not one of the EAI\_\* values, a pointer to a string indicating an unknown error is returned.

### Related Information

The getaddrinfo Subroutine, freeaddrinfo Subroutine, and getnameinfo Subroutine articles in *AIX 5L Version 5.3 Technical Reference: Communications Volume 2*.

**Subroutines Overview** in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## gamma Subroutine

### Purpose

Computes the natural logarithm of the gamma function.

### Libraries

The **gamma**:  
IEEE Math Library (**libm.a**)  
or System V Math Library (**libmsaa.a**)



## Syntax

```
#include <math.h>
extern int signgam;
double gamma (x)
double x;
```

## Description

The **gamma** subroutine computes the logarithm of the gamma function.

The sign of gamma( *x* ) is returned in the external integer **signgam**.

**Note:** Compile any routine that uses subroutines from the **libm.a** with the **-lm** flag. To compile the **lgamma.c** file, enter:

```
cc lgamma.c -lm
```

## Parameters

*x* Specifies the value to be computed.

## Related Information

“exp, expf, or expl Subroutine” on page 245, “feclearexcept Subroutine” on page 265, “fetetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

The **exp**, **expm1**, **log**, **log10**, **log1p** or **pow** (“exp, expf, or expl Subroutine” on page 245) subroutine, **matherr** (“matherr Subroutine” on page 789) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## gencore or coredump Subroutine

### Purpose

Creates a core file without terminating the process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <core.h>

int gencore (coredumpinfo)
struct coredumpinfo *coredumpinfo;

int coredump (coredumpinfo)
struct coredumpinfo *coredumpinfo;
```

## Description

The **gencore** and **coredump** subroutines create a core file of a process without terminating it. The core file contains the snapshot of the process at the time the call is made and can be used with the **dbx** command for debugging purposes.

If any thread of the process is in a system call when its snapshot core file is generated, the register information returned may not be reliable (except for the stack pointer). To save all user register contents when a system call is made so that they are available to the **gencore** and **coredump** subroutines, the application should be built using the "-bM:UR" flags.

If any thread of the process is sleeping inside the kernel or stopped (possibly for job control), the caller of the **gencore** and **coredump** subroutines will also be blocked until the thread becomes runnable again. Thus, these subroutines may take a long time to complete depending upon the target process state.

The **coredump** subroutine always generates a core file for the process from which it is called. This subroutine has been replaced by the **gencore** subroutine and is being provided for compatibility reasons only.

The **gencore** subroutine creates a core file for the process whose process ID is specified in the *pid* field of the **coredumpinfo** structure. For security measures, the user ID (uid) and group ID (gid) of the core file are set to the uid and gid of the process.

Both these subroutines return success even if the core file cannot be created completely because of filesystem space constraints. When using the **dbx** command with an incomplete core file, **dbx** may warn that the core file is truncated.

In the "Change / Show Characteristics of Operating System" smitty screen, there are two options regarding the creation of the core file. The core file will always be created in the default core format and will ignore the value specified in the "Use pre-430 style CORE dump" option. However, the value specified for the "Enable full CORE dump" option will be considered when creating the core file. Resource limits of the target process for **file** and **coredump** will be enforced.

The **coredumpinfo** structure contains the following fields:

Member Type	Member Name	Description
unsigned int	<i>length</i>	Length of the core file name.
char *	<i>name</i>	Name of the core file.
pid_t	<i>pid</i>	ID of the process to be coredumped.
int	<i>flags</i>	Flags-version flag. Set this to <b>COREGEN_VERSION_1</b> .

**Note:** The *pid* and *flags* fields are required for the **gencore** subroutine, but are ignored for the **coredump** subroutine

## Parameters

*coredumpinfo* Specifies the address of the **coredumpinfo** structure that provides the file name to save the core snapshot and its length. For the **gencore** subroutine, it also provides the process id of the process whose core is to be dumped and a flag which includes version flag bits. The version flag value must be set to **COREGEN\_VERSION\_1**.

## Return Values

Upon successful completion, the **gencore** and **coredump** subroutines return a 0. If unsuccessful, a -1 is returned, and the **errno** global variable is set to indicate the error

## Error Codes

<b>EACCES</b>	Search permission is denied on a component of the path prefix, the file exists and permissions specified by the mode are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.
<b>ENOENT</b>	The <i>name</i> field in the <i>coredumpinfo</i> parameter points to an empty string.
<b>EINTR</b>	The subroutine was interrupted by a signal before it could complete.
<b>ENAMETOOLONG</b>	The value of the <i>length</i> field in the <b>coredumpinfo</b> structure or the length of the absolute path of the specified core file name is greater than <b>MAXPATHLEN</b> (as defined in the <b>sys/param.h</b> file).
<b>EINVAL</b>	The value of the <i>length</i> field in the <b>coredumpinfo</b> structure is 0.
<b>EAGAIN</b>	The target process is already in the middle of another <b>gencore</b> or <b>coredump</b> subroutine.
<b>ENOMEM</b>	Unable to allocate memory resources to complete the subroutine.

In addition to the above, the following **errno** values can be set when the **gencore** subroutine is unsuccessful:

<b>EPERM</b>	The real or effective user ID of the calling process does not match the real or effective user ID of target process or the calling process does not have root user authority.
<b>ESRCH</b>	There is no process whose ID matches the value specified in the <i>pid</i> field of the <i>coredumpinfo</i> parameter or the process is exiting.
<b>EINVAL</b>	The <i>flags</i> field in the <i>coredumpinfo</i> parameter is not set to a valid version value.

## Related Information

The adb Command, in *AIX 5L Version 5.3 Commands Reference, Volume 1*.

The dbx command, and gencore Command in *AIX 5L Version 5.3 Commands Reference, Volume 2*.

The core file format in *AIX 5L Version 5.3 Files Reference*.

---

## genpagvalue Subroutine

### Purpose

Sets the current process credentials.

### Library

Security Library (libc.a)

### Syntax

```
#include <pag.h>
int genpagvalue(pag_name, pag_value, pag_flags);
char *      pag_name;
uint64_t *  pag_value;
int         pag_flags;
```

## Description

The **genpagvalue** subroutine generates a new PAG value for a given PAG name. For this function to succeed, the PAG name must be registered with the operating system before calling the **genpagvalue** subroutine. The **genpagvalue** subroutine is limited to maintaining information about the last generated PAG number and accordingly generating a new number. This service can optionally store the PAG value in the process's **cred** structure. It does not monitor the PAG values stored in the **cred** structure by other means.

The PAG value returned is of size 64 bits. The number of significant bits is determined by the requested PAG type. 32-bit PAGs have 32 significant bits. 64-bit PAGs have 62 significant bits.

A process must have root authority to invoke this function for 32-bit PAG types. Any process may invoke this function for 64-bit PAG types.

The *pag\_flags* parameter with the value **PAG\_SET\_VALUE** causes the generated value to be atomically stored in the process's credentials. The *pag\_flags* parameter with both the **PAG\_SET\_VALUE** and **PAG\_COPY\_CRED** values set causes the current process's credentials to be duplicated before the generated value is stored.

## Parameters

<i>pag_name</i>	The name parameter is a 1 to 4 character, NULL terminated name for the PAG type. Typical values include afs, dfs, pki and krb5.
<i>pag_value</i>	This pointer points to a buffer where the OS will return the newly generated PAG value.
<i>pag_flags</i>	These flags control the behavior of the <b>getpagvalue</b> subroutine. This must be set to 0 or one or more of the values <b>PAG_SET_VALUE</b> or <b>PAG_COPY_CRED</b> .

## Return Values

A value of 0 is returned upon successful completion. If the **genpagvalue** subroutine fails a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **genpagvalue** subroutine fails if one or more of the following are true:

<b>EINVAL</b>	The PAG value cannot be generated because the named PAG type does not exist as part of the table.
<b>EPERM</b>	The process does not have the correct authority to use the service.

Other errors might be set by subroutines invoked by the **genpagvalue** subroutine.

## Related Information

`__pag_getid` System Call, `__pag_getname` System Call, `__pag_getvalue` System Call, `__pag_setname` System Call, `__pag_setvalue` System Call, `kcred_genpagvalue` Kernel Service, `kcred_getpagid` Kernel Service, `kcred_getpagname` Kernel Service in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 1*.

List of Security and Auditing Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## get\_malloc\_log Subroutine

### Purpose

Retrieves information about the malloc subsystem.

### Syntax

```
#include <malloc.h>
size_t get_malloc_log (addr, buf, bufsize)
void *addr;
void *buf;
size_t bufsize;
```

### Description

The **get\_malloc\_log** subroutine retrieves a record of currently active malloc allocations. These records are stored as an array of **malloc\_log** structures, which are copied from the process heap into the buffer specified by the *buf* parameter. No more than *bufsize* bytes are copied into the buffer. Only records corresponding to the heap of which *addr* is a member are copied, unless *addr* is NULL, in which case records from all heaps are copied. The *addr* parameter must be either a pointer to space allocated previously by the malloc subsystem or NULL.

### Parameters

<i>addr</i>	Pointer to a space allocated by the malloc subsystem.
<i>buf</i>	Specifies into which buffer the <b>malloc_log</b> structures are stored.
<i>bufsize</i>	Specifies the number of bytes that can be copied into the buffer.

### Return Values

The **get\_malloc\_log** subroutine returns the number of bytes actually transferred into the *bufsize* parameter. If Malloc Log is not enabled, 0 is returned. If *addr* is not a pointer allocated by the malloc subsystem, 0 is returned and the **errno** global variable is set to **EINVAL**.

### Related Information

“malloc, free, realloc, calloc, mallopt, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778, and “get\_malloc\_log\_live Subroutine.”

reset\_malloc\_log Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## get\_malloc\_log\_live Subroutine

### Purpose

Provides information about the malloc subsystem.

### Syntax

```
#include <malloc.h>
struct malloc_log* get_malloc_log_live (addr)
void *addr;
```

## Description

The **get\_malloc\_log\_live** subroutine provides access to a record of currently active malloc allocations. The information is stored as an array of **malloc\_log** structures, which are located in the process heap. This data is volatile and subject to update. The *addr* parameter must be either a pointer to space allocated previously by the malloc subsystem or NULL.

## Parameters

*addr*                      Pointer to space allocated previously by the malloc subsystem

## Return Values

The **get\_malloc\_log\_live** subroutine returns a pointer to the process heap at which the records of current malloc allocations are stored. If the *addr* parameter is NULL, a pointer to the beginning of the array is returned. If *addr* is a pointer to space allocated previously by the malloc subsystem, the pointer returned corresponds to records of the same heap as *addr*. If Malloc Log is not enabled, NULL is returned. If *addr* is not a pointer allocated by the malloc subsystem, NULL is returned and the **errno** global variable is set to **EINVAL**.

## Related Information

“malloc, free, realloc, calloc, mallopt, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778, and “get\_malloc\_log Subroutine” on page 341.

reset\_malloc\_log Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## get\_speed, set\_speed, or reset\_speed Subroutines

### Purpose

Set and get the terminal baud rate.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/str_tty.h>
int get_speed (FileDescriptor)
int FileDescriptor;
int set_speed (FileDescriptor, Speed)
int FileDescriptor;
int Speed;
int reset_speed (FileDescriptor)
int FileDescriptor;
```

### Description

The baud rate functions **set\_speed** subroutine and **get\_speed** subroutine are provided to allow the user applications to program any value of the baud rate that is supported by the asynchronous adapter, but that cannot be expressed using the termios subroutines **cfsetospeed**, **cfsetispeed**, **cfgetospeed**, and **cfgetispeed**. Those subroutines are indeed limited to the set values {BO, B50, ..., B38400} described in **<termios.h>**.

### Interaction with the termios Baud flags:

If the terminal's device driver supports these subroutines, it has two interfaces for baud rate manipulation.

### Operation for Baud Rate:

normal mode: This is the default mode, in which a termios supported speed is in use.

speed-extended mode: This mode is entered either by calling **set\_speed** subroutine a non-termios supported speed at the configuration of the line.

In this mode, all the calls to **tcgetattr** subroutine or **TCGETS ioctl** subroutine will have B50 in the returned termios structure.

If **tcsetatt** subroutine or **TCSETS**, **TCSETAF**, or **TCSETAW ioctl** subroutines is called and attempt to set B50, the actual baud rate is not changed. If it attempts to set any other termios-supported speed, the driver will switch back to the normal mode and the requested baud rate is set. Calling **reset\_speed** subroutine is another way to switch back to the normal mode.

## Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>Speed</i>	The integer value of the requested speed.

## Return Values

Upon successful completion, **set\_speed** and **reset\_speed** return a value of 0, and **get\_speed** returns a positive integer specifying the current speed of the line. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

**EINVAL** The *FileDescriptor* parameter does not specify a valid file descriptor for a **tty** the recognizes the **set\_speed**, **get\_speed** and **reset\_speed** subroutines, or the *Speed* parameter of **set\_speed** is not supported by the terminal.

Plus all the **errno** codes that may be set in case of failure in an **ioctl** subroutine issued to a streams based **tty**.

## Related Information

**cfgetospeed**, **cfsetospeed**, **cfgetispeed**, or **cfsetispeed** (“cfgetospeed, cfsetospeed, cfgetispeed, or cfsetispeed Subroutine” on page 142) subroutines.

---

## getargs Subroutine

### Purpose

Gets arguments of a process.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <procinfo.h>
#include <sys/types.h>
```

```

int getargs (processBuffer, bufferLen, argsBuffer, argsLen)
struct procsinfo *processBuffer
or struct procentry64 *processBuffer;
int bufferLen;
char *argsBuffer;
int argsLen;

```

## Description

The **getargs** subroutine returns a list of parameters that were passed to a command when it was started. Only one process can be examined per call to **getargs**.

The **getargs** subroutine uses the `pi_pid` field of *processBuffer* to determine which process to look for. *bufferLen* should be set to the size of **struct procsinfo** or **struct procentry64**. Parameters are returned in *argsBuffer*, which should be allocated by the caller. The size of this array must be given in *argsLen*.

On return, *argsBuffer* consists of a succession of strings, each terminated with a null character (ascii `\0`). Hence, two consecutive NULLs indicate the end of the list.

**Note:** The arguments may be changed asynchronously by the process, but results are not guaranteed to be consistent.

## Parameters

*processBuffer*

Specifies the address of a **procsinfo** or **procentry64** structure, whose `pi_pid` field should contain the pid of the process that is to be looked for.

*bufferLen*

Specifies the size of a single **procsinfo** or **procentry64** structure.

*argsBuffer*

Specifies the address of an array of characters to be filled with a series of strings representing the parameters that are needed. An extra NULL character marks the end of the list. This array must be allocated by the caller.

*argsLen*

Specifies the size of the *argsBuffer* array. No more than *argsLen* characters are returned.

## Return Values

If successful, the **getargs** subroutine returns zero. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **getargs** subroutine does not succeed if the following are true:

<b>ESRCH</b>	The specified process does not exist.
<b>EFAULT</b>	The copy operation to the buffer was not successful or the <i>processBuffer</i> or <i>argsBuffer</i> parameters are invalid.
<b>EINVAL</b>	The <i>bufferLen</i> parameter does not contain the size of a single <b>procsinfo</b> or <b>procentry64</b> structure.
<b>ENOMEM</b>	There is no memory available in the address space.



## Related Information

The **getevvars** (“getevvars Subroutine” on page 365), **getpid** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getpgrp** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getppid** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getprocs** or **getthrds** (“getthrds Subroutine” on page 443) subroutines.

The **ps** command.

---

## getaudithostattr, IDtohost, hosttoID, nexthost or putaudithostattr Subroutine

### Purpose

Accesses the host information in the audit host database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int getaudithostattr (Hostname, Attribute, Value, Type)
char *Hostname;
char *Attribute;
void *Value;
int Type;
```

```
char *IDtohost (ID);
char *ID;
```

```
char *hosttoID (Hostname, Count);
char *Hostname;
int Count;
```

```
char *nexthost (void);
```

```
int putaudithostattr (Hostname, Attribute, Value, Type);
char *Hostname;
char *Attribute;
void *Value;
int Type;
```

### Description

These subroutines access the audit host information.

The **getaudithostattr** subroutine reads a specified attribute from the host database. If the database is not already open, this subroutine does an implicit open for reading.

Similarly the **putaudithostattr** subroutine writes a specified attribute into the host database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by the **putaudithostattr** must be explicitly committed by calling the **putaudithostattr** subroutine with a Type of **SEC\_COMMIT**. Until all the data is committed, only these subroutines within the process return written data.

New entries in the host database must first be created by invoking **putaudithostattr** with the **SEC\_NEW** type.

The **IDtohost** subroutine converts an 8 byte host identifier into a hostname.

The **hosttoid** subroutine converts a hostname to a pointer to an array of valid 8 byte host identifiers. A pointer to the array of identifiers is returned on success. A **NULL** pointer is returned on failure. The number of known host identifiers is returned in **\*Count**.

The **nexthost** subroutine returns a pointer to the name of the next host in the audit host database.

## Parameters

<i>Attribute</i>	Specifies which attribute is read. The following possible attributes are defined in the <b>usersec.h</b> file:  <b>S_AUD_CPUID</b> Host identifier list. The attribute type is <b>SEC_LIST</b> .
<i>Count</i>	Specifies the number of 8 byte host identifier entries that are available in the <i>IDarray</i> parameter or that have been returned in the <i>IDarray</i> parameter.
<i>Hostname</i>	Specifies the name of the host for the operation.
<i>ID</i>	An 8 byte host identifier.
<i>IDarray</i>	Specifies a pointer to an array of 1 or more 8 byte host identifiers.
<i>Type</i>	Specifies the type of attribute expected. Valid types are defined in <b>usersec.h</b> . The only valid Type value is <b>SEC_LIST</b> .
<i>Value</i>	The return value for read operations and the new value for write operations.

## Return Values

On successful completion, the **getaudithostattr**, **IDtohost**, **hosttoid**, **nexthost**, or **putaudithostattr** subroutine returns 0. If unsuccessful, the subroutine returns non-zero.

## Error Codes

The **getaudithostattr**, **IDtohost**, **hosttoid**, **nexthost**, or **putaudithostattr** subroutine fails if the following is true:

<b>EINVAL</b>	If invalid attribute <i>Name</i> or if <i>Count</i> is equal to zero for the <b>hosttoid</b> subroutine.
<b>ENOENT</b>	If there is no matching <i>Hostname</i> entry in the database.

## Related Information

The **auditmerge** command, **auditpr** command, **auditselect** command, **auditstream** command.

The **auditread** (“auditread, auditread\_r Subroutines” on page 111) subroutine.

---

## getauthdb or getauthdb\_r Subroutine

### Purpose

Finds the current administrative domain.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <usersec.h>
```

```
int getauthdb (Value)  
authdb_t *Value;
```

```
int getauthdb_r (Value)  
authdb_t *Value;
```

## Description

The **getauthdb** and **getauthdb\_r** subroutines return the value of the current authentication domain in the *Value* parameter. The **getauthdb** subroutine returns the value of the current process-wide authentication domain. The **getauthdb\_r** subroutine returns the authentication domain for the current thread if one has been set. The subroutines return -1 if no administrative domain has been set.

## Parameters

<i>Value</i>	A pointer to a variable of type <b>authdb_t</b> . The <b>authdb_t</b> type is a 16-character array that contains the name of a loadable authentication module.
--------------	--

## Return Values

1	The value returned is from the process-wide data.
0	The value returned is from the thread-specific data. An authentication database module has been specified by an earlier call to the <b>setauthdb</b> subroutine. The name of the current database module has been copied to the <i>Value</i> parameter.
-1	The subroutine failed. An authentication database module has not been specified by an earlier call to the <b>setauthdb</b> subroutine.

## Related Information

setauthdb or setauthdb\_r Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## getc, getchar, fgetc, or getw Subroutine

### Purpose

Gets a character or word from an input stream.

### Library

Standard I/O Package (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
int getc ( Stream)  
FILE *Stream;
```

```
int fgetc (Stream)  
FILE *Stream;
```

```
int getchar (void)
int getw (Stream)
FILE *Stream;
```

## Description

The **getc** macro returns the next byte as an **unsigned char** data type converted to an **int** data type from the input specified by the *Stream* parameter and moves the file pointer, if defined, ahead one byte in the *Stream* parameter. The **getc** macro cannot be used where a subroutine is necessary; for example, a subroutine pointer cannot point to it.

Because it is implemented as a macro, the **getc** macro does not work correctly with a *Stream* parameter value that has side effects. In particular, the following does not work:

```
getc(*f++)
```

In such cases, use the **fgetc** subroutine.

The **fgetc** subroutine performs the same function as the **getc** macro, but **fgetc** is a true subroutine, not a macro. The **fgetc** subroutine runs more slowly than **getc** but takes less disk space.

The **getchar** macro returns the next byte from **stdin** (the standard input stream). The **getchar** macro is equivalent to **getc(stdin)**.

The first successful run of the **fgetc**, **fgets**, **fgetc**, **fgetwc**, **fgetws**, **fread**, **fscanf**, **getc**, **getchar**, **gets** or **scanf** subroutine using a stream that returns data not supplied by a prior call to the **ungetc** or **ungetwc** subroutine marks the *st\_atime* field for update.

The **getc** and **getchar** macros have also been implemented as subroutines for ANSI compatibility. To access the subroutines instead of the macros, insert **#undef getc** or **#undef getchar** at the beginning of the source file.

The **getw** subroutine returns the next word (**int**) from the input specified by the *Stream* parameter and increments the associated file pointer, if defined, to point to the next word. The size of a word varies from one machine architecture to another. The **getw** subroutine returns the constant **EOF** at the end of the file or when an error occurs. Since **EOF** is a valid integer value, the **feof** and **ferror** subroutines should be used to check the success of **getw**. The **getw** subroutine assumes no special alignment in the file.

Because of additional differences in word length and byte ordering from one machine architecture to another, files written using the **putw** subroutine are machine-dependent and may not be readable using the **getw** macro on a different type of processor.

## Parameters

*Stream*            Points to the file structure of an open file.

## Return Values

Upon successful completion, the **getc**, **fgetc**, **getchar**, and **getw** subroutines return the next byte or **int** data type from the input stream pointed by the *Stream* parameter. If the stream is at the end of the file, an end-of-file indicator is set for the stream and the integer constant **EOF** is returned. If a read error occurs, the **errno** global variable is set to reflect the error, and a value of **EOF** is returned. The **ferror** and **feof** subroutines should be used to distinguish between the end of the file and an error condition.

## Error Codes

If the stream specified by the *Stream* parameter is unbuffered or data needs to be read into the stream's buffer, the **getc**, **getchar**, **fgetc**, or **getw** subroutine is unsuccessful under the following error conditions:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> flag is set for the file descriptor underlying the stream specified by the <i>Stream</i> parameter. The process would be delayed in the <b>fgetc</b> subroutine operation.
<b>EBADF</b>	Indicates that the file descriptor underlying the stream specified by the <i>Stream</i> parameter is not a valid file descriptor opened for reading.
<b>EFBIG</b>	Indicates that an attempt was made to read a file that exceeds the process' file-size limit or the maximum file size. See the <b>ulimit</b> subroutine.
<b>EINTR</b>	Indicates that the read operation was terminated due to the receipt of a signal, and either no data was transferred, or the implementation does not report partial transfer for this file. <b>Note:</b> Depending upon which library routine the application binds to, this subroutine may return <b>EINTR</b> . Refer to the <b>signal</b> subroutine regarding <b>sa_restart</b> .
<b>EIO</b>	Indicates that a physical error has occurred, or the process is in a background process group attempting to perform a <b>read</b> subroutine call from its controlling terminal, and either the process is ignoring (or blocking) the <b>SIGTTIN</b> signal or the process group is orphaned.
<b>EPIPE</b>	Indicates that an attempt is made to read from a pipe or first-in-first-out (FIFO) that is not open for reading by any process. A <b>SIGPIPE</b> signal will also be sent to the process.
<b>EOVERFLOW</b>	Indicates that the file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding stream.

The **getc**, **getchar**, **fgetc**, or **getw** subroutine is also unsuccessful under the following error conditions:

<b>ENOMEM</b>	Indicates insufficient storage space is available.
<b>ENXIO</b>	Indicates either a request was made of a nonexistent device or the request was outside the capabilities of the device.

## Related Information

The **feof**, **ferror**, **clearerr**, or **fileno** (“feof, ferror, clearerr, or fileno Macro” on page 270) subroutine, **freopen**, **fopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fread** or **fwrite** (“fread or fwrite Subroutine” on page 310) subroutine, **getwc**, **fgetwc**, or **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **get** or **fgets** (“gets or fgets Subroutine” on page 434) subroutine, **putc**, **putchar**, **fputc**, or **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **scanf**, **sscanf**, **fscanf**, or **wscanf** subroutine.

List of Character Manipulation Services, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getc\_unlocked, getchar\_unlocked, putc\_unlocked, putchar\_unlocked Subroutines

### Purpose

stdio with explicit client locking.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
int getc_unlocked (FILE * stream);
int getchar_unlocked (void);
int putc_unlocked (int c, FILE * stream);
int putchar_unlocked (int c);
```

## Description

Versions of the functions **getc**, **getchar**, **putc**, and **putchar** respectively named **getc\_unlocked**, **getchar\_unlocked**, **putc\_unlocked**, and **putchar\_unlocked** are provided which are functionally identical to the original versions with the exception that they are not required to be implemented in a thread-safe manner. They may only safely be used within a scope protected by **flockfile** (or **ftrylockfile**) and **funlockfile**. These functions may safely be used in a multi-threaded program if and only if they are called while the invoking thread owns the (FILE\*) object, as is the case after a successful call of the **flockfile** or **ftrylockfile** functions.

## Return Values

See **getc**, **getchar**, **putc**, and **putchar**.

## Related Information

The **getc** or **getchar** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **putc** or **putchar** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine.

---

## getconfattr or putconfattr Subroutine

### Purpose

Accesses the system information in the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
#include <userconf.h>
```

```
int getconfattr (sys, Attribute, Value, Type)
char * sys;
char * Attribute;
void *Value;
int Type;
```

```
int putconfattr (sys, Attribute, Value, Type)
char * sys;
char * Attribute;
void *Value;
int Type;
```

### Description

The **getconfattr** subroutine reads a specified attribute from the system information database. The **putconfattr** subroutine writes a specified attribute to the system information database.

### Parameters

**sys** System attribute. The following possible attributes are defined in the **userconf.h** file.

- SC\_SYS\_LOGIN
- SC\_SYS\_USER
- SC\_SYS\_ADMUSER
- SC\_SYS\_AUDIT     SEC\_LIST
- SC\_SYS\_AUSERS    SEC\_LIST
- SC\_SYS\_ASYS      SEC\_LIST
- SC\_SYS\_ABIN      SEC\_LIST
- SC\_SYS\_ASTREAM   SEC\_LIST

Users can define parameter the system attribute parameter. In this case the parameter value is used as a stanza name. The stanza name contains the specified attribute and value in the *Attribute* and *Value* parameters. The **putconfattr** subroutine creates this stanza in the file associated with the attribute. The **getconfattr** subroutine retrieves the value for the specified attribute and the user defined stanza.

#### *Attribute*

Specifies which attribute is read. The following possible attributes are defined in the **usersec.h** file:

#### **S\_CORECOMP**

Core compression status. The attribute type is **SEC\_CHAR**.

#### **S\_COREPATH**

Core path specification status. The attribute type is **SEC\_CHAR**.

#### **S\_COREPNAME**

Core path specification location. The attribute type is **SEC\_CHAR**.

#### **S\_CORENAMING**

Core naming status. The attribute type is **SEC\_CHAR**.

#### **S\_PGRP**

Principle group name. The attribute type is **SEC\_CHAR**.

#### **S\_GROUPS**

Groups to which the user belongs. The attribute type is **SEC\_LIST**.

#### **S\_ADMGROUPS**

Groups for which the user is an administrator. The attribute type is **SEC\_LIST**.

#### **S\_ADMIN**

Administrative status of a user. The attribute type is **SEC\_BOOL**.

#### **S\_AUDITCLASSES**

Audit classes to which the user belongs. The attribute type is **SEC\_LIST**.

#### **S\_AUTHSYSTEM**

Defines the user's authentication method. The attribute type is **SEC\_CHAR**.

#### **S\_HOME**

Home directory. The attribute type is **SEC\_CHAR**.

#### **S\_SHELL**

Initial program run by a user. The attribute type is **SEC\_CHAR**.

#### **S\_GECOS**

Personal information for a user. The attribute type is **SEC\_CHAR**.

#### **S\_USRENV**

User-state environment variables. The attribute type is **SEC\_LIST**.

#### **S\_SYSENV**

Protected-state environment variables. The attribute type is **SEC\_LIST**.

**S\_LOGINCHK**

Specifies whether the user account can be used for local logins. The attribute type is **SEC\_BOOL**.

**S\_HISTEXPIRE**

Defines the period of time (in weeks) that a user cannot reuse a password. The attribute type is **SEC\_INT**.

**S\_HISTSIZE**

Specifies the number of previous passwords that the user cannot reuse. The attribute type is **SEC\_INT**.

**S\_MAXREPEAT**

Defines the maximum number of times a user can repeat a character in a new password. The attribute type is **SEC\_INT**.

**S\_MINAGE**

Defines the minimum age in weeks that the user's password must exist before the user can change it. The attribute type is **SEC\_INT**.

**S\_PWDCHECKS**

Defines the password restriction methods for this account. The attribute type is **SEC\_LIST**.

**S\_MINALPHA**

Defines the minimum number of alphabetic characters required in a new user's password. The attribute type is **SEC\_INT**.

**S\_MINDIFF**

Defines the minimum number of characters required in a new password that were not in the old password. The attribute type is **SEC\_INT**.

**S\_MINLEN**

Defines the minimum length of a user's password. The attribute type is **SEC\_INT**.

**S\_MINOTHER**

Defines the minimum number of non-alphabetic characters required in a new user's password. The attribute type is **SEC\_INT**.

**S\_DICTIIONLIST**

Defines the password dictionaries for this account. The attribute type is **SEC\_LIST**.

**S\_SUCHK**

Specifies whether the user account can be accessed with the **su** command. Type **SEC\_BOOL**.

**S\_REGISTRY**

Defines the user's authentication registry. The attribute type is **SEC\_CHAR**.

**S\_RLOGINCHK**

Specifies whether the user account can be used for remote logins using the **telnet** or **rlogin** commands. The attribute type is **SEC\_BOOL**.

**S\_DAEMONCHK**

Specifies whether the user account can be used for daemon execution of programs and subsystems using the **cron** daemon or **src**. The attribute type is **SEC\_BOOL**.

**S\_TPATH**

Defines how the account may be used on the trusted path. The attribute type is **SEC\_CHAR**. This attribute must be one of the following values:

**nosak** The secure attention key is not enabled for this account.

**notsh** The trusted shell cannot be accessed from this account.



**always**

This account may only run trusted programs.

**on**

Normal trusted-path processing applies.

**S\_TTYS**

List of ttys that can or cannot be used to access this account. The attribute type is **SEC\_LIST**.

**S\_SUGROUPS**

Groups that can or cannot access this account. The attribute type is **SEC\_LIST**.

**S\_EXPIRATION**

Expiration date for this account, in seconds since the epoch. The attribute type is **SEC\_CHAR**.

**S\_AUTH1**

Primary authentication methods for this account. The attribute type is **SEC\_LIST**.

**S\_AUTH2**

Secondary authentication methods for this account. The attribute type is **SEC\_LIST**.

**S\_UFSIZE**

Process file size soft limit. The attribute type is **SEC\_INT**.

**S\_UCPU**

Process CPU time soft limit. The attribute type is **SEC\_INT**.

**S\_UDATA**

Process data segment size soft limit. The attribute type is **SEC\_INT**.

**S\_USTACK**

Process stack segment size soft limit. Type: **SEC\_INT**.

**S\_URSS**

Process real memory size soft limit. Type: **SEC\_INT**.

**S\_UCORE**

Process core file size soft limit. The attribute type is **SEC\_INT**.

**S\_PWD**

Specifies the value of the passwd field in the **/etc/passwd** file. The attribute type is **SEC\_CHAR**.

**S\_UMASK**

File creation mask for a user. The attribute type is **SEC\_INT**.

**S\_LOCKED**

Specifies whether the user's account can be logged into. The attribute type is **SEC\_BOOL**.

**S\_UFSIZE\_HARD**

Process file size hard limit. The attribute type is **SEC\_INT**.

**S\_UCPU\_HARD**

Process CPU time hard limit. The attribute type is **SEC\_INT**.

**S\_UDATA\_HARD**

Process data segment size hard limit. The attribute type is **SEC\_INT**.

**S\_USTACK\_HARD**

Process stack segment size hard limit. Type: **SEC\_INT**.

### **S\_URSS\_HARD**

Process real memory size hard limit. Type: **SEC\_INT**.

### **S\_UCORE\_HARD**

Process core file size hard limit. The attribute type is **SEC\_INT**.

**Note:** These values are string constants that should be used by applications both for convenience and to permit optimization in latter implementations.

*Type* Specifies the type of attribute expected. Valid types are defined in the **usersec.h** file and include:

#### **SEC\_INT**

The format of the attribute is an integer.

For the **getconfattr** subroutine, the user should supply a pointer to a defined integer variable. For the **putconfattr** subroutine, the user should supply an integer.

#### **SEC\_CHAR**

The format of the attribute is a null-terminated character string.

#### **SEC\_LIST**

The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series is terminated by two successive null characters.

#### **SEC\_BOOL**

The format of the attribute from the **getconfattr** subroutine is an integer with the value of either 0 (false) or 1 (true). The format of the attribute for the **putconfattr** subroutine is a null-terminated string containing one of the following strings: true, false, yes, no, always, or never.

#### **SEC\_COMMIT**

For the **putconfattr** subroutine, this value specified by itself indicates that the changes to the named *sys* value or stanza are to be committed to permanent storage. The *Attribute* and *Value* parameters are ignored. If no stanza name is specified, all outstanding changes to the system information databases are committed to permanent storage.

#### **SEC\_DELETE**

The corresponding attribute is deleted from the database.

## **Security**

Files Accessed:

<b>Mode</b>	<b>File</b>
<b>rw</b>	/etc/security/user
<b>rw</b>	/etc/security/limits
<b>rw</b>	/etc/security/login.cfg
<b>rw</b>	/usr/lib/security/mkuser.default
<b>rw</b>	/etc/security/audit/config

## **Return Values**

If successful, the **getconfattr** subroutine returns a value of zero.

If unsuccessful, the **getconfattr** subroutine returns a value of -1.

## Error Codes

<b>ENOENT</b>	The value that the <i>Sys</i> parameter specifies does not exist.
<b>ENOATTR</b>	The specified <i>Attribute</i> variable is not defined for this <i>Sys</i> parameter.
<b>EINVAL</b>	The <i>Attribute</i> or <i>Type</i> variable for the specified <i>Sys</i> parameter is not valid.
<b>EACCESS</b>	The user does not have access to the specified <i>Attribute</i> variable.

## Files

`/etc/passwd`                      Contains user IDs.

## Related Information

The “getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getconfattrs Subroutine

### Purpose

Accesses system information in the system information database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
#include <userconf.h>
```

```
int getconfattrs (Sys, Attributes, Count)
char * Sys;
dbattr_t * Attributes;
int Count
```

### Description

The **getconfattrs** subroutine accesses system configuration information.

The **getconfattrs** subroutine reads one or more attributes from the system configuration database. If the database is not already open, this subroutine does an implicit open for reading.

The *Attributes* array contains information about each attribute that is to be written. The **dbattr\_t** data structure contains the following fields:

**attr\_name**  
The name of the desired attribute.

**attr\_idx**  
Used internally by the **getconfattrs** subroutine.

**attr\_type**  
The type of the desired attribute. The list of attribute types is defined in the **usersec.h** header file.

### **attr\_flag**

The results of the request to read the desired attribute.

### **attr\_un**

A union containing the values to be written. Its union members that follow correspond to the definitions of the **attr\_char**, **attr\_int**, **attr\_long**, and **attr\_llong** macros, respectively:

#### **un\_char**

Attributes of type **SEC\_CHAR** and **SEC\_LIST** store a pointer to the value to be written.

**un\_int** Attributes of type **SEC\_INT** and **SEC\_BOOL** contain the value of the attribute to be written.

#### **un\_long**

Attributes of type **SEC\_LONG** contain the value of the attribute to be written.

#### **un\_llong**

Attributes of type **SEC\_LLONG** contain the value of the attribute to be written.

### **attr\_domain**

The authentication domain containing the attribute. The **getconfattr** subroutine is responsible for managing the memory referenced by this pointer.

Use the **setuserdb** and **enduserdb** subroutines to open and close the system configuration database. Failure to explicitly open and close the system database can result in loss of memory and performance.

## **Parameters**

<i>Sys</i>	Specifies the name of the subsystem for which the attributes are to be read.
<i>Attributes</i>	A pointer to an array of one or more elements of type <b>dbattr_t</b> . The list of system attributes is defined in the <b>usersec.h</b> header file.
<i>Count</i>	The number of array elements in <i>Attributes</i> .

## **Security**

Files accessed:

<b>Mode</b>	<b>File</b>
r	<b>/etc/security/.ids</b>
r	<b>/etc/security/audit/config</b>
r	<b>/etc/security/audit/events</b>
r	<b>/etc/security/audit/objects</b>
r	<b>/etc/security/login.cfg</b>
r	<b>/etc/security/portlog</b>
r	<b>/etc/security/roles</b>
r	<b>/usr/lib/security/methods.cfg</b>
r	<b>/usr/lib/security/mkuser.default</b>

## **Return Values**

If the value of the *Sys* or *Attributes* parameter is NULL, or the value of the *Count* parameter is less than 1, the **getconfattr** subroutine returns a value of -1, and sets the **errno** global variable to indicate the error. Otherwise, the subroutine returns a value of zero. The **getconfattr** subroutine does not check the validity of the *Sys* parameter. Each element in the *Attributes* array must be examined on a successful call to the **getconfattr** subroutine to determine whether the *Attributes* array entry is successfully retrieved.

## Error Codes

The **getconfattr** subroutine returns an error that indicates that the system attribute does or does not exist. Additional errors can indicate an error querying the information databases for the requested attributes.

<b>EINVAL</b>	The <i>Attributes</i> parameter is NULL.
<b>EINVAL</b>	The <i>Count</i> parameter is less than 1.
<b>ENOENT</b>	The specified <i>Sys</i> does not exist.

If the **getconfattr** subroutine fails to query an attribute, one or more of the following errors is returned in the **attr\_flag** field of the corresponding *Attributes* element:

<b>EACCES</b>	The user does not have access to the attribute specified in the <b>attr_name</b> field.
<b>EINVAL</b>	The <b>attr_type</b> field in the <i>Attributes</i> entry contains an invalid type.
<b>EINVAL</b>	The <b>attr_un</b> field in the <i>Attributes</i> entry does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors might not be detected.
<b>ENOMEM</b>	Memory could not be allocated to store the return value.
<b>ENOATTR</b>	The <b>attr_name</b> field in the <i>Attributes</i> entry specifies an attribute that is not defined for this system table.

## Files

<b>/etc/security/ids</b>	The next available user and group ID values.
<b>/etc/security/audit/config</b>	Bin and stream mode audit configuration parameters.
<b>/etc/security/audit/events</b>	Format strings for audit event records.
<b>/etc/security/audit/objects</b>	File system objects that are explicitly audited.
<b>/etc/security/login.cfg</b>	Miscellaneous login relation parameters.
<b>/etc/security/portlog</b>	Port login failure and locking history.
<b>/etc/security/roles</b>	Definitions of administrative roles.
<b>/usr/lib/security/methods.cfg</b>	Definitions of loadable authentication modules.
<b>/usr/lib/security/mkuser.default</b>	Default user attributes for administrative and nonadministrative users.

## Related Information

The “getconfattr or putconfattr Subroutine” on page 350.

The “getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454.

The “getconfattr or putconfattr Subroutine” on page 350.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getcontext or setcontext Subroutine

### Purpose

Initializes the structure pointed to by *ucp* to the context of the calling process.

### Library

(libc.a)

## Syntax

```
#include <ucontext.h>
```

```
int getcontext (ucontext_t *ucp);
```

```
int setcontext (const ucontext_t *ucp);
```

## Description

The **getcontext** subroutine initializes the structure pointed to by *ucp* to the current user context of the calling process. The **ucontext\_t** type that *ucp* points to defines the user context and includes the contents of the calling process' machine registers, the signal mask, and the current execution stack.

The **setcontext** subroutine restores the user context pointed to by *ucp*. A successful call to **setcontext** subroutine does not return; program execution resumes at the point specified by the *ucp* argument passed to **setcontext** subroutine. The *ucp* argument should be created either by a prior call to **getcontext** subroutine, or by being passed as an argument to a signal handler. If the *ucp* argument was created with **getcontext** subroutine, program execution continues as if the corresponding call of **getcontext** subroutine had just returned. If the *ucp* argument was created with **makecontext** subroutine, program execution continues with the function passed to **makecontext** subroutine. When that function returns, the process continues as if after a call to **setcontext** subroutine with the *ucp* argument that was input to **makecontext** subroutine. If the *ucp* argument was passed to a signal handler, program execution continues with the program instruction following the instruction interrupted by the signal. If the *uc\_link* member of the **ucontext\_t** structure pointed to by the *ucp* argument is equal to 0, then this context is the main context, and the process will exit when this context returns.

## Parameters

*ucp*                      A pointer to a user structure.

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **getcontext** and **setcontext** subroutines are unsuccessful if one of the following is true:

<b>EINVAL</b>	NULL <i>ucp</i> address
<b>EFAULT</b>	Invalid <i>ucp</i> address

## Related Information

The **makecontext** ("makecontext or swapcontext Subroutine" on page 788) subroutine, **setjmp** subroutine, **sigaltstack** subroutine, **sigaction** subroutine, **sigprocmask** subroutine, and **sigsetjmp** subroutine.

---

## getcwd Subroutine

### Purpose

Gets the path name of the current directory.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <unistd.h>
```

```
char *getcwd ( Buffer, Size)  
char *Buffer;  
size_t Size;
```

## Description

The **getcwd** subroutine places the absolute path name of the current working directory in the array pointed to by the *Buffer* parameter, and returns that path name. The *size* parameter specifies the size in bytes of the character array pointed to by the *Buffer* parameter.

## Parameters

*Buffer* Points to string space that will contain the path name. If the *Buffer* parameter value is a null pointer, the **getcwd** subroutine, using the **malloc** subroutine, obtains the number of bytes of free space as specified by the *Size* parameter. In this case, the pointer returned by the **getcwd** subroutine can be used as the parameter in a subsequent call to the **free** subroutine. Starting the **getcwd** subroutine with a null pointer as the *Buffer* parameter value is not recommended.

*Size* Specifies the length of the string space. The value of the *Size* parameter must be at least 1 greater than the length of the path name to be returned.

## Return Values

If the **getcwd** subroutine is unsuccessful, a null value is returned and the **errno** global variable is set to indicate the error. The **getcwd** subroutine is unsuccessful if the *Size* parameter is not large enough or if an error occurs in a lower-level function.

In UNIX03 mode, the **getcwd** subroutine returns a null value if the actual path name is longer than the value defined by **PATH\_MAX** (see the **limits.h** file). In the previous mode, the **getcwd** subroutine returns a truncated path name if the path name is longer than **PATH\_MAX**. The previous behavior is disabled by setting the environment variable **XPG\_SUS\_ENV=ON**.

## Error Codes

If the **getcwd** subroutine is unsuccessful, it returns one or more of the following error codes:

<b>EACCES</b>	Indicates that read or search permission was denied for a component of the path name
<b>EINVAL</b>	Indicates that the <i>Size</i> parameter is 0 or a negative number.
<b>ENOMEM</b>	Indicates that insufficient storage space is available.
<b>ERANGE</b>	Indicates that the <i>Size</i> parameter is greater than 0, but is smaller than the length of the path name plus 1.

## Related Information

The “getwd Subroutine” on page 479, “malloc, free, realloc, calloc, malloc, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getdate Subroutine

### Purpose

Convert user format date and time.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <time.h>
```

```
struct tm *getdate (const char *string)
extern int getdate_err
```

### Description

The **getdate** subroutine converts user definable date and/or time specifications pointed to by *string*, into a **struct tm**. The structure declaration is in the **time.h** header file (see **ctime** subroutine).

User supplied templates are used to parse and interpret the input string. The templates are contained in text files created by the user and identified by the environment variable **DATEMSK**. The **DATEMSK** variable should be set to indicate the full pathname of the file that contains the templates. The first line in the template that matches the input specification is used for interpretation and conversion into the internal time format.

The templates should follow a format where complex field descriptors are preceded by simpler ones. For example:

```
%M
%H:%M
%m/%d/%y
%m/%d/%y %H:%M:%S
```

The following field descriptors are supported:

<b>%%</b>	Same as %.
<b>%a</b>	Abbreviated weekday name.
<b>%A</b>	Full weekday name.
<b>%b</b>	Abbreviated month name.
<b>%B</b>	Full month name.
<b>%c</b>	Locale's appropriate date and time representation.
<b>%C</b>	Century number (00-99; leading zeros are permitted but not required)
<b>%d</b>	Day of month (01 - 31: the leading zero is optional).
<b>%e</b>	Same as %d.
<b>%D</b>	Date as %m/%d/%y.
<b>%h</b>	Abbreviated month name.
<b>%H</b>	Hour (00 - 23)
<b>%I</b>	Hour (01 - 12)
<b>%m</b>	Month number (01 - 12)
<b>%M</b>	Minute (00 - 59)
<b>%n</b>	Same as \n.
<b>%p</b>	Locale's equivalent of either AM or PM.
<b>%r</b>	Time as %I:%M:%S %p
<b>%R</b>	Time as %H: %M
<b>%S</b>	Seconds (00 - 61) Leap seconds are allowed but are not predictable through use of algorithms.
<b>%t</b>	Same as tab.



<b>%T</b>	Time as %H: %M: %S
<b>%w</b>	Weekday number (Sunday = 0 - 6)
<b>%x</b>	Locale's appropriate date representation.
<b>%X</b>	Locale's appropriate time representation.
<b>%y</b>	Year within century. <b>Note:</b> When the environment variable <b>XPG_TIME_FMT=ON</b> , <b>%y</b> is the year within the century. When a century is not otherwise specified, values in the range 69-99 refer to years in the twentieth century (1969 to 1999, inclusive); values in the range 00-68 refer to 2000 to 2068, inclusive.
<b>%Y</b>	Year as ccyy (such as 1986)
<b>%Z</b>	Time zone name or no characters if no time zone exists. If the time zone supplied by <b>%Z</b> is not the same as the time zone <b>getdate</b> subroutine expects, an invalid input specification error will result. The <b>getdate</b> subroutine calculates an expected time zone based on information supplied to the interface (such as hour, day, and month).

The match between the template and input specification performed by the **getdate** subroutine is case sensitive.

The month and weekday names can consist of any combination of upper and lower case letters. The user can request that the input date or time specification be in a specific language by setting the **LC\_TIME** category (See the **setlocale** subroutine).

Leading zero's are not necessary for the descriptors that allow leading zero's. However, at most two digits are allowed for those descriptors, including leading zero's. Extra whitespace in either the template file or in *string* is ignored.

The field descriptors **%c**, **%x**, and **%X** will not be supported if they include unsupported field descriptors.

Example 1 is an example of a template. Example 2 contains valid input specifications for the template. Example 3 shows how local date and time specifications can be defined in the template.

The following rules apply for converting the input specification into the internal format:

- If only the weekday is given, today is assumed if the given month is equal to the current day and next week if it is less.
- If only the month is given, the current month is assumed if the given month is equal to the current month and next year if it is less and no year is given (the first day of month is assumed if no day is given).
- If no hour, minute, and second are given, the current hour, minute and second are assumed.
- If no date is given, today is assumed if the given hour is greater than the current hour and tomorrow is assumed if it is less.

## Return Values

Upon successful completion, the **getdate** subroutine returns a pointer to **struct tm**; otherwise, it returns a null pointer and the external variable **getdate\_err** is set to indicate the error.

## Error Codes

Upon failure, a null pointer is returned and the variable **getdate\_err** is set to indicate the error.

The following is a complete list of the **getdate\_err** settings and their corresponding descriptions:

- |          |  |
|----------|--|
| <b>1</b> | The <b>DATMSK</b> environment variable is null or undefined. |
| <b>2</b> | The template file cannot be opened for reading.              |
| <b>3</b> | Failed to get file status information.                       |

- 4 The template file is not a regular file.
- 5 An error is encountered while reading the template file.
- 6 Memory allocation failed (not enough memory available).
- 7 There is no line in the template that matches the input.
- 8 Invalid input specification, Example: February 31 or a time is specified that can not be represented in a time\_t (representing the time in seconds since 00:00:00 UTC, January 1, 1970).

## Examples

- The following example shows the possible contents of a template:

```
%m
%A %B %d, %Y, %H:%M:%S
%A
%B
%m/%d/%y %I %p
%d, %m, %Y %H:%M
at %A the %dst of %B in %Y
run job at %I %p, %B %dnd
&A den %d. %B %Y %H.%M Uhr
```

- The following are examples of valid input specifications for the template in Example 1:

```
getdate ("10/1/87 4 PM")
getdate ("Friday")
getdate ("Friday September 18, 1987, 10:30:30")
getdate ("24,9,1986 10:30")
getdate ("at monday the 1st of december in 1986")
getdate ("run job at 3 PM. december 2nd")
```

If the LC\_TIME category is set to a German locale that includes freitag as a weekday name and oktober as a month name, the following would be valid:

```
getdate ("freitag den 10. oktober 1986 10.30 Uhr")
```

- The following examples shows how local date and time specification can be defined in the template.

Invocation	Line in Template
getdate ("11/27/86")	%m/%d/%y
getdate ("27.11.86"0	%d.%m.%y
getdate ("86-11-27")	%y-%m-%d
getdate ("Friday 12:00:00")	%A %H:%M:%S

- The following examples help to illustrate the above rules assuming that the current date Mon Sep 22 12:19:47 EDT 1986 and the LC\_TIME category is set to the default "C" locale.

Input	Line in Template	Date
Mon	%a	Mon Sep 22 12:19:47 EDT 1986
Sun	%a	Sun Sep 28 12:19:47 EDT 1986
Fri	%a	Fri Sep 26 12:19:47 EDT 1986
September	%B	Mon Sep1 12:19:47 EDT 1986
January	%B	Thu Jan 1 12:19:47 EDT 1986
December	%B	Mon Dec 1 12:19:47 EDT 1986
Sep Mon	%b %a	Mon Sep 1 12:19:47 EDT 1986
Jan Fri	%b %a	Fri Jan 2 12:19:47 EDT 1986
Dec Mon	%b %a	Mon Dec 1 12:19:47 EDT 1986
Jan Wed 1989	%b %a %Y	Wed Jan 4 12:19:47 EDT 1986

Input	Line in Template	Date
Fri 9	%a %H	Fri Sep 26 12:19:47 EDT 1986
Feb 10:30	%b %H: %S	Sun Feb 1 12:19:47 EDT 1986
10:30	%H: %M	Tue Sep 23 12:19:47 EDT 1986
13:30	%H: %M	Mon Sep 22 12:19:47 EDT 1986

## Related Information

The **ctime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200), **ctype** (“ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines” on page 208), **setlocale**, **strftime**, and **times** (“getrusage, getrusage64, times, or vtimes Subroutine” on page 428) subroutines.

---

## getdtablesize Subroutine

### Purpose

Gets the descriptor table size.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int getdtablesize (void)
```

### Description

The **getdtablesize** subroutine is used to determine the size of the file descriptor table.

The size of the file descriptor table for a process is set by the **ulimit** command or by the **setrlimit** subroutine. The **getdtablesize** subroutine returns the current size of the table as reported by the **getrlimit** subroutine. If **getrlimit** reports that the table size is unlimited, **getdtablesize** instead returns the value of **OPEN\_MAX**, which is the largest possible size of the table.

**Note:** The **getdtablesize** subroutine returns a runtime value that is specific to the version of the operating system on which the application is running. In AIX 4.3.1, **getdtablesize** returns a value that is set in the **limits** file, which can be different from system to system.

### Return Values

The **getdtablesize** subroutine returns the size of the descriptor table.

## Related Information

The **close** (“close Subroutine” on page 175) subroutine, **open** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **select** subroutine.

---

## getea Subroutine

### Purpose

Reads the value of an extended attribute.

## Syntax

```
#include <sys/ea.h>
```

```
ssize_t getea(const char *path, const char *name,  
              void *value, size_t size);  
ssize_t fgetea(int filedes, const char *name, void *value, size_t size);  
ssize_t lgetea(const char *path, const char *name,  
               void *value, size_t size);
```

## Description

Extended attributes are name:value pairs associated with the file system objects (such as files, directories, and symlinks). They are extensions to the normal attributes that are associated with all objects in the file system (that is, the **stat(2)** data).

Do not define an extended attribute name with the eight characters prefix "(0xF8)SYSTEM(0xF8)". Prefix "(0xF8)SYSTEM(0xF8)" is reserved for system use only.

**Note:** The 0xF8 prefix represents a non-printable character.

The **getea** subroutine retrieves the value of the extended attribute identified by *name* and associated with the given *path* in the file system. The length of the attribute *value* is returned. The **fgetea** subroutine is identical to **getea**, except that it takes a file descriptor instead of a path. The **lgetea** subroutine is identical to **getea**, except, in the case of a symbolic link, the link itself is interrogated rather than the file that it refers to.

## Parameters

<i>path</i>	The path name of the file.
<i>name</i>	The name of the extended attribute. An extended attribute name is a NULL-terminated string.
<i>value</i>	A pointer to a buffer in which the attribute will be stored. The value of an extended attribute is an opaque byte stream of specified length.
<i>size</i>	The size of the buffer. If <i>size</i> is 0, <b>getea</b> returns the current size of the named extended attribute, which can be used to estimate whether the size of a buffer is sufficiently large enough to hold the value associated with the extended attribute.
<i>filedes</i>	A file descriptor for the file.

## Return Values

If the **getea** subroutine succeeds, a nonnegative number is returned that indicates the size of the extended attribute value. Upon failure, -1 is returned and **errno** is set appropriately.

## Error Codes

<b>EACCES</b>	Caller lacks read permission on the base file, or lacks the appropriate ACL privileges for named attribute <b>read</b> .
<b>EFAULT</b>	A bad address was passed for <i>path</i> , <i>name</i> , or <i>value</i> .
<b>EFORMAT</b>	File system is capable of supporting EAs, but EAs are disabled.
<b>EINVAL</b>	A path-like name should not be used (such as <b>zml/file</b> , <b>.</b> and <b>..</b> ).
<b>ENAMETOOLONG</b>	The <i>path</i> or <i>name</i> value is too long.
<b>ENOATTR</b>	The named attribute does not exist, or the process has no access to this attribute.
<b>ERANGE</b>	The size of the value buffer is too small to hold the result.
<b>ENOTSUP</b>	Extended attributes are not supported by the file system.

The errors documented for the **stat(2)** system call are also applicable here.

## Related Information

“listea Subroutine” on page 724, removeea Subroutine, setea Subroutine, stateea Subroutine.

---

## getenv Subroutine

### Purpose

Returns the value of an environment variable.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
char *getenv ( Name )  
const char *Name;
```

### Description

The **getenv** subroutine searches the environment list for a string of the form *Name=Value*. Environment variables are sometimes called shell variables because they are frequently set with shell commands.

### Parameters

*Name*            Specifies the name of an environment variable. If a string of the proper form is not present in the current environment, the **getenv** subroutine returns a null pointer.

### Return Values

The **getenv** subroutine returns a pointer to the value in the current environment, if such a string is present. If such a string is not present, a null pointer is returned. The **getenv** subroutine normally does not modify the returned string. The **putenv** subroutine, however, may overwrite or change the returned string. Do not attempt to free the returned pointer. The **getenv** subroutine returns a pointer to the user’s copy of the environment (which is static), until the first invocation of the **putenv** subroutine that adds a new environment variable. The **putenv** subroutine allocates an area of memory large enough to hold both the user’s environment and the new variable. The next call to the **getenv** subroutine returns a pointer to this newly allocated space that is not static. Subsequent calls by the **putenv** subroutine use the **realloc** subroutine to make space for new variables. Unsuccessful completion returns a null pointer.

## Related Information

The **putenv** (“putenv Subroutine” on page 1326) subroutine.

---

## getenvvars Subroutine

### Purpose

Gets environment of a process.

### Library

Standard C library (**libc.a**)

## Syntax

```
#include <procinfo.h>
#include <sys/types.h>
```

```
int getevars (processBuffer, bufferLen, argsBuffer, argsLen)
struct procsinfo *processBuffer
or struct procsinfo64 *processBuffer;
int bufferLen;
char *argsBuffer;
int argsLen;
```

## Description

The **getevars** subroutine returns the environment that was passed to a command when it was started. Only one process can be examined per call to **getevars**.

The **getevars** subroutine uses the *pi\_pid* field of *processBuffer* to determine which process to look for. *bufferLen* should be set to size of **struct procsinfo** or **struct procsinfo64**. Parameters are returned in *argsBuffer*, which should be allocated by the caller. The size of this array must be given in *argsLen*.

On return, *argsBuffer* consists of a succession of strings, each terminated with a null character (ascii '\0'). Hence, two consecutive NULLs indicate the end of the list.

**Note:** The arguments may be changed asynchronously by the process, but results are not guaranteed to be consistent.

## Parameters

*processBuffer*

Specifies the address of a **procsinfo** or **procsinfo64** structure, whose *pi\_pid* field should contain the pid of the process that is to be looked for.

*bufferLen*

Specifies the size of a single **procsinfo** or **procsinfo64** structure.

*argsBuffer*

Specifies the address of an array of characters to be filled with a series of strings representing the parameters that are needed. An extra NULL character marks the end of the list. This array must be allocated by the caller.

*argsLen*

Specifies the size of the *argsBuffer* array. No more than *argsLen* characters are returned.

## Return Values

If successful, the **getevars** subroutine returns zero. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **getevars** subroutine does not succeed if the following are true:

<b>ESRCH</b>	The specified process does not exist.
<b>EFAULT</b>	The copy operation to the buffer was not successful or the <i>processBuffer</i> or <i>argsBuffer</i> parameters are invalid.
<b>EINVAL</b>	The <i>bufferLen</i> parameter does not contain the size of a single <b>procsinfo</b> or <b>procsinfo64</b> structure.
<b>ENOMEM</b>	There is no memory available in the address space.

## Related Information

The **getargs** (“getargs Subroutine” on page 343), **getpid** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getpgrp** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getppid** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getprocs** or **getthrds** (“getthrds Subroutine” on page 443) subroutines.

The **ps** command.

---

## getfilehdr Subroutine

### Purpose

Retrieves the header details of the advanced accounting data file.

### Library

The **libaacct.a** library.

### Syntax

```
#define <sys/aacct.h>
getfilehdr(filename, hdrinfo)
char *filename;
struct aacct_file_header *hdrinfo;
```

### Description

The **getfilehdr** subroutine retrieves the advanced accounting data file header information in a structure of type **aacct\_file\_header** and returns it to the caller through the structure pointer passed to it. The data file header contains the system details such as the name of the host, the partition number, and the system model.

### Parameters

<i>filename</i>	Name of the advanced accounting data file.
<i>hdrinfo</i>	Pointer to the <b>aacct_file_header</b> structure in which the header information is returned.

### Security

No restrictions. Any user can call this function.

### Return Values

0	The call to the subroutine was successful.
-1	The call to the subroutine failed.

### Error Codes

<b>EINVAL</b>	The passed pointer is NULL.
<b>ENOENT</b>	Specified data file does not exist.
<b>EPERM</b>	Permission denied. Unable to read the data file.

## Related Information

The “buildproclist Subroutine” on page 125, “buildtranlist or freetranlist Subroutine” on page 126, “getproclist, getlparlist, or getarmlist Subroutine” on page 414.

## getfirstprojdb Subroutine

### Purpose

Retrieves details of the first project from the specified project database.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
getfirstprojdb(void *handle, struct project *project, char *comm)
```

### Description

The **getfirstprojdb** subroutine retrieves the first project definitions from the project database, which is controlled through the *handle* parameter. The caller must initialize the project database prior to calling this routine with the **projdballoc** routine. Upon successful completion, the project information is copied to the project structure specified by the caller. In addition, the associated project comment, if present, is copied to the buffer pointed to by the *comm* parameter. The comment buffer is allocated by the caller and must have a length of 1024 bytes.

There is an internal state (that is, the current project) associated with the project database. When the project database is initialized, the current project is the first project in the database. The **getnextprojdb** subroutine returns the current project and advances the current project assignment to the next project in the database so that successive calls read each project entry in the database. The **getfirstprojdb** subroutine can be used to reset the database, so that the initial project is the current project assignment.

### Parameters

<i>handle</i>	Pointer to the <b>projdb</b> handle.
<i>project</i>	Pointer to project structure where the retrieved data is stored.
<i>comm</i>	Pointer to the comment buffer.

### Security

No restriction. Any user can call this function.

### Return Values

0	Success
-1	Failure

### Error Codes

<b>EINVAL</b>	Invalid arguments, if passed pointer is NULL.
<b>ENOENT</b>	No projects available.



## Related Information

The “addprojdb Subroutine” on page 32, “chprojattrdb Subroutine” on page 159, “getnextprojdb Subroutine” on page 396, “getprojdb Subroutine” on page 419, “getprojs Subroutine” on page 420, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, “projdbfree Subroutine” on page 1181, rprojdb Subroutine.

---

## getfsent, getfsspec, getfsfile, getfstype, setfsent, or endfsent Subroutine

### Purpose

Gets information about a file system.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <fstab.h>
struct fstab *getfsent( )

struct fstab *getfsspec ( Special)
char *Special;

struct fstab *getfsfile( File)
char *File;

struct fstab *getfstype( Type)
char *Type;
void setfsent( )
void endfsent( )
```

### Description

The **getfsent** subroutine reads the next line of the **/etc/filesystems** file, opening the file if necessary.

The **setfsent** subroutine opens the **/etc/filesystems** file and positions to the first record.

The **endfsent** subroutine closes the **/etc/filesystems** file.

The **getfsspec** and **getfsfile** subroutines sequentially search from the beginning of the file until a matching special file name or file-system file name is found, or until the end of the file is encountered. The **getfstype** subroutine does likewise, matching on the file-system type field.

**Note:** All information is contained in a static area, which must be copied to be saved.

### Parameters

<i>Special</i>	Specifies the file-system file name.
<i>File</i>	Specifies the file name.
<i>Type</i>	Specifies the file-system type.

## Return Values

The **getfsent**, **getfsspec**, **getfstype**, and **getfsfile** subroutines return a pointer to a structure that contains information about a file system. The header file **fstab.h** describes the structure. A null pointer is returned when the end of file (EOF) is reached or if an error occurs.

## Files

**/etc/filesystems** Centralizes file system characteristics.

## Related Information

The **getvfsent**, **getvfsbytype**, **getvfsbyname**, **getvfsbyflag**, **setvfsent**, or **endvfsent** (“getvfsent, getvfsbytype, getvfsbyname, getvfsbyflag, setvfsent, or endvfsent Subroutine” on page 476) subroutine.

The **filesystems** file.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getgid, getegid or gegidx Subroutine

### Purpose

Gets the process group IDs.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
#include <sys/types.h>
gid_t getgid (void);
gid_t getegid (void);
#include <id.h>
gid_t gegidx (int type);
```

### Description

The **getgid** subroutine returns the real group ID of the calling process.

The **getegid** subroutine returns the effective group ID of the calling process.

The **gegidx** subroutine returns the group ID indicated by the *type* parameter of the calling process.

These subroutines are part of Base Operating System (BOS) Runtime.

### Return Values

The **getgid**, **getegid** and **gegidx** subroutines return the requested group ID. The **getgid** and **getegid** subroutines are always successful.

The **gegidx** subroutine will return -1 and set the global **errno** variable to **EINVAL** if the *type* parameter is not one of **ID\_REAL**, **ID\_EFFECTIVE** or **ID\_SAVED**.

## Parameters

*type* Specifies the group ID to get. Must be one of **ID\_REAL** (real group ID), **ID\_EFFECTIVE** (effective group ID) or **ID\_SAVED** (saved set-group ID).

## Error Codes

If the **getgidx** subroutine fails the following is returned:

**EINVAL** Indicates the value of the type parameter is invalid.

## Related Information

The “getgroups Subroutine” on page 383, **initgroups** subroutine, **setgid** subroutine, **setgroups** subroutine.

The **groups** command, **setgroups** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine

### Purpose

Accesses the basic group information in the user database.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <grp.h>
struct group *getgrent ( );

struct group *getgrgid (GID)
gid_t GID;

struct group *getgrnam (Name)
const char * Name;
void setgrent ( );
void endgrent ( );
```

### Description

**Attention:** The information returned by the **getgrent**, **getgrnam**, and **getgrgid** subroutines is stored in a static area and is overwritten on subsequent calls. You must copy this information to save it.

**Attention:** These subroutines should not be used with the **getgroupattr** subroutine. The results are unpredictable.

The **setgrent** subroutine opens the user database if it is not already open. Then, this subroutine sets the cursor to point to the first group entry in the database.

The **getgrent**, **getgrnam**, and **getgrgid** subroutines return information about the requested group. The **getgrent** subroutine returns the next group in the sequential search. The **getgrnam** subroutine returns the first group in the database whose name matches that of the *Name* parameter. The **getgrgid** subroutine returns the first group in the database whose group ID matches the *GID* parameter. The **endgrent** subroutine closes the user database.

**Note:** An ! (exclamation mark) is written into the *gr\_passwd* field. This field is ignored and is present only for compatibility with older versions of UNIX.

These subroutines also return the list of user members for the group. Currently, the list is limited to 2000 entries (this could change in the future to where all the entries for the group are returned).

## The Group Structure

The **group** structure is defined in the **grp.h** file and has the following fields:

<i>gr_name</i>	Contains the name of the group.
<i>gr_passwd</i>	Contains the password of the group. <b>Note:</b> This field is no longer used.
<i>gr_gid</i>	Contains the ID of the group.
<i>gr_mem</i>	Contains the members of the group.

If the Network Information Service (NIS) is enabled on the system, these subroutines attempt to retrieve the group information from the NIS authentication server.

## Parameters

<i>GID</i>	Specifies the group ID.
<i>Name</i>	Specifies the group name.

*Group* Specifies the basic group information to enter into the user database.

## Return Values

If successful, the **getgrent**, **getgrnam**, and **getgrgid** subroutines return a pointer to a valid group structure. Otherwise, a null pointer is returned.

## Error Codes

These subroutines fail if one or more of the following are returned:

<b>EIO</b>	Indicates that an input/output (I/O) error has occurred.
<b>EINTR</b>	Indicates that a signal was caught during the <b>getgrnam</b> or <b>getgrgid</b> subroutine.
<b>EMFILE</b>	Indicates that the maximum number of file descriptors specified by the <b>OPEN_MAX</b> value are currently open in the calling process.
<b>ENFILE</b>	Indicates that the maximum allowable number of files is currently open in the system.

To check an application for error situations, set the **errno** global variable to a value of 0 before calling the **getgrgid** subroutine. If the **errno** global variable is set on return, an error occurred.

## File

*/etc/group* Contains basic group attributes.

## Related Information

“putgrent Subroutine” on page 1326

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getgrgid\_r Subroutine

### Purpose

Gets a group database entry for a group ID.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
#include <sys/types.h>
#include <grp.h>

int getgrgid_r(gid_t gid,
struct group *grp,
char *buffer,
size_t bufsize,
struct group **result);
```

### Description

The **getgrgid\_r** subroutine updates the **group** structure pointed to by *grp* and stores a pointer to that structure at the location pointed to by *result*. The structure contains an entry from the group database with a matching *gid*. Storage referenced by the group structure is allocated from the memory provided with the *buffer* parameter, which is *bufsize* characters in size. The maximum size needed for this buffer can be determined with the `{_SC_GETGR_R_SIZE_MAX}` *sysconf* parameter. A NULL pointer is returned at the location pointed to by *result* on error or if the requested entry is not found.

### Return Values

Upon successful completion, **getgrgid\_r** returns a pointer to a **struct group** with the structure defined in `<grp.h>` with a matching entry if one is found. The **getgrgid\_r** function returns a null pointer if either the requested entry was not found, or an error occurred. On error, *errno* will be set to indicate the error.

The return value points to a static area that is overwritten by a subsequent call to the **getgrent**, **getgrgid**, or **getgrnam** subroutine.

If successful, the **getgrgid\_r** function returns zero. Otherwise, an error number is returned to indicate the error.

### Error Codes

The **getgrgid\_r** function fails if:

<b>ERANGE</b>	Insufficient storage was supplied via <i>buffer</i> and <i>bufsize</i> to contain the data to be referenced by the resulting <b>group</b> structure.
---------------	--

Applications wishing to check for error situations should set *errno* to 0 before calling **getgrgid\_r**. If *errno* is set on return, an error occurred.

## Related Information

The **getgrent**, **getgrgid**, **getgrnam**, **setgrent**, **endgrent** (“getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371) subroutine.

The `<grp.h>`, `<limits.h>`, and `<sys/types.h>` header files.

---

## getgrnam\_r Subroutine

### Purpose

Search a group database for a name.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
#include <sys/types.h>
#include <grp.h>

int getgrnam_r (const char **name,
                struct group *grp,
                char *buffer,
                size_t bufsize,
                struct group **result);
```

### Description

The **getgrnam\_r** function updates the **group** structure pointed to by *grp* and stores pointer to that structure at the location pointed to by *result*. The structure contains an entry from the group database with a matching *gid* or *name*. Storage referenced by the group structure is allocated from the memory provided with the *buffer* parameter, which is *bufsize* characters in size. The maximum size needed for this buffer can be determined with the `{_SC_GETGR_R_SIZE_MAX}` *sysconf* parameter. A NULL pointer is returned at the location pointed to by *result* on error or if the requested entry is not found.

### Return Values

The **getgrnam\_r** function returns a pointer to a **struct group** with the structure defined in `<grp.h>` with a matching entry if one is found. The **getgrnam\_r** function returns a null pointer if either the requested entry was not found, or an error occurred. On error, *errno* will be set to indicate the error.

The return value points to a static area that is overwritten by a subsequent call to the **getgrent**, **getgrgid**, or **getgrnam** subroutine.

If successful, the **getgrnam\_r** function returns zero. Otherwise, an error number is returned to indicate the error.

### Error Codes

The **getgrnam\_r** function fails if:

<b>ERANGE</b>	Insufficient storage was supplied via <i>buffer</i> and <i>bufsize</i> to contain the data to be referenced by the resulting <b>group</b> structure.
---------------	--

Applications wishing to check for error situations should set *errno* to 0 before calling **getgrnam\_r**. If *errno* is set on return, an error occurred.

## Related Information

The **getgrent**, **getgrgid**, **getgrnam**, **setgrent**, **endgrent** (“getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371) subroutine.

The `<grp.h>`, `<limits.h>`, and `<sys/types.h>` header files.

---

## getgroupattr, IDtogroup, nextgroup, or putgroupattr Subroutine

### Purpose

Accesses the group information in the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int getgroupattr (Group, Attribute, Value, Type)
```

```
char * Group;
```

```
char * Attribute;
```

```
void * Value;
```

```
int Type;
```

```
int putgroupattr (Group, Attribute, Value, Type)
```

```
char *Group;
```

```
char *Attribute;
```

```
void *Value;
```

```
int Type;
```

```
char *IDtogroup ( GID)
```

```
gid_t GID;
```

```
char *nextgroup ( Mode, Argument)
```

```
int Mode, Argument;
```

### Description

**Attention:** These subroutines and the **setpwent** and **setgrent** subroutines should not be used simultaneously. The results can be unpredictable.

These subroutines access group information. Because of their greater granularity and extensibility, you should use them instead of the **getgrent**, **putgrent**, **getgrnam**, **getgrgid**, **setgrent**, and **endgrent** subroutines.

The **getgroupattr** subroutine reads a specified attribute from the group database. If the database is not already open, the subroutine will do an implicit open for reading.

Similarly, the **putgroupattr** subroutine writes a specified attribute into the group database. If the database is not already open, the subroutine does an implicit open for reading and writing. Data changed by **putgroupattr** must be explicitly committed by calling the **putgroupattr** subroutine with a *Type* parameter specifying the **SEC\_COMMIT** value. Until the data is committed, only **get** subroutine calls within the process will return the written data.

New entries in the user and group databases must first be created by invoking **putgroupattr** with the **SEC\_NEW** type.

The **IDtogroup** subroutine translates a group ID into a group name.

The **nextgroup** subroutine returns the next group in a linear search of the group database. The consistency of consecutive searches depends upon the underlying storage-access mechanism and is not guaranteed by this subroutine.

The **setuserdb** and **enduserdb** subroutines should be used to open and close the user database.

## Parameters

*Argument*

Presently unused and must be specified as null.

*Attribute*

Specifies which attribute is read. The following possible values are defined in the **usersec.h** file:

**S\_ID** Group ID. The attribute type is **SEC\_INT**.

**S\_USERS**

Members of the group. The attribute type is **SEC\_LIST**.

**S\_ADMS**

Administrators of the group. The attribute type is **SEC\_LIST**.

**S\_ADMIN**

Administrative status of a group. Type: **SEC\_BOOL**.

**S\_GRPEXPORT**

Specifies if the DCE registry can overwrite the local group information with the DCE group information during a DCE export operation. The attribute type is **SEC\_BOOL**.

Additional user-defined attributes may be used and will be stored in the format specified by the *Type* parameter.

*GID*

Specifies the group ID to be translated into a group name.

*Group*

Specifies the name of the group for which an attribute is to be read.

*Mode*

Specifies the search mode. Also can be used to delimit the search to one or more user credential databases. Specifying a non-null *Mode* value implicitly rewinds the search. A null mode continues the search sequentially through the database. This parameter specifies one of the following values as a bit mask (defined in the **usersec.h** file):

**S\_LOCAL**

The local database of groups are included in the search.

**S\_SYSTEM**

All credentials servers for the system are searched.



<i>Type</i>	<p>Specifies the type of attribute expected. Valid values are defined in the <b>usersec.h</b> file and include:</p> <p><b>SEC_INT</b> The format of the attribute is an integer. The buffer returned by the <b>getgroupattr</b> subroutine and the buffer supplied by the <b>putgroupattr</b> subroutine are defined to contain an integer.</p> <p><b>SEC_CHAR</b> The format of the attribute is a null-terminated character string.</p> <p><b>SEC_LIST</b> The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series is terminated by two successive null characters.</p> <p><b>SEC_BOOL</b> A pointer to an integer (<b>int *</b>) that has been cast to a null pointer.</p> <p><b>SEC_COMMIT</b> For the <b>putgroupattr</b> subroutine, this value specified by itself indicates that changes to the named group are committed to permanent storage. The <i>Attribute</i> and <i>Value</i> parameters are ignored. If no group is specified, changes to all modified groups are committed to permanent storage.</p> <p><b>SEC_DELETE</b> The corresponding attribute is deleted from the database.</p> <p><b>SEC_NEW</b> If using the <b>putgroupattr</b> subroutine, updates all the group database files with the new group name.</p>
<i>Value</i>	<p>Specifies the address of a pointer for the <b>getgroupattr</b> subroutine. The <b>getgroupattr</b> subroutine will return the address of a buffer in the pointer. For the <b>putgroupattr</b> subroutine, the <i>Value</i> parameter specifies the address of a buffer in which the attribute is stored. See the <i>Type</i> parameter for more details.</p>

## Security

Files Accessed:

Mode	File
rw	<b>/etc/group</b> (write access for <b>putgroupattr</b> )
rw	<b>/etc/security/group</b> (write access for <b>putgroupattr</b> )

## Return Values

The **getgroupattr** and **putgroupattr** subroutines, when successfully completed, return a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

The **IDtgroup** and **nextgroup** subroutines return a character pointer to a buffer containing the requested group name, if successfully completed. Otherwise, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

**Note:** All of these subroutines return errors from other subroutines.

These subroutines fail if the following is true:

<b>EACCES</b>	Access permission is denied for the data request.
---------------	---

The **getgroupattr** and **putgroupattr** subroutines fail if one or more of the following are true:

<b>EINVAL</b>	The <i>Value</i> parameter does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors may not be detected.
<b>EINVAL</b>	The <i>Type</i> parameter contains more than one of the <b>SEC_INT</b> , <b>SEC_BOOL</b> , <b>SEC_CHAR</b> , <b>SEC_LIST</b> , or <b>SEC_COMMIT</b> attributes.
<b>EINVAL</b>	The <i>Type</i> parameter specifies that an individual attribute is to be committed, and the <i>Group</i> parameter is null.
<b>ENOENT</b>	The specified <i>Group</i> parameter does not exist or the attribute is not defined for this group.
<b>EPERM</b>	Operation is not permitted.

The **IDtogroup** subroutine fails if the following is true:

<b>ENOENT</b>	The <i>GID</i> parameter could not be translated into a valid group name on the system.
---------------	---

The **nextgroup** subroutine fails if one or more of the following are true:

<b>EINVAL</b>	The <i>Mode</i> parameter is not null, and does not specify either <b>S_LOCAL</b> or <b>S_SYSTEM</b> .
<b>EINVAL</b>	The <i>Argument</i> parameter is not null.
<b>ENOENT</b>	The end of the search was reached.

## Related Information

The **getuserattr** (“getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454) subroutine, **getuserpw** (“getuserpw, putuserpw, or putuserpwhist Subroutine” on page 468) subroutine, **setpwdb** subroutine, **setuserdb** subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getgroupattrs Subroutine

### Purpose

Retrieves multiple group attributes in the group database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int getgroupattrs (Group, Attributes, Count)
char * Group;
dbattr_t * Attributes;
int Count
```

### Description

**Attention:** Do not use this subroutine and the **setpwent** and **setgrent** subroutines simultaneously. The results can be unpredictable.

The **getgroupattrs** subroutine accesses group information. Because of its greater granularity and extensibility, use it instead of the **getgrent** routines.

The **getgroupattrs** subroutine reads one or more attributes from the group database. If the database is not already open, this subroutine does an implicit open for reading. A call to the **getgroupattrs** subroutine with an *Attributes* parameter of null and *Count* parameter of 0 for every new group verifies that the group exists.

The *Attributes* array contains information about each attribute that is to be read. The **dbattr\_t** data structure contains the following fields:

**attr\_name**

The name of the desired attribute.

**attr\_idx**

Used internally by the **getgroupattrs** subroutine.

**attr\_type**

The type of the desired attribute. The list of attribute types is defined in the **usersec.h** header file.

**attr\_flag**

The results of the request to read the desired attribute.

**attr\_un**

A union containing the returned values. Its union members that follow correspond to the definitions of the **attr\_char**, **attr\_int**, **attr\_long**, and **attr\_llong** macros, respectively:

**un\_char**

Attributes of type **SEC\_CHAR** and **SEC\_LIST** store a pointer to the returned attribute in this member when the requested attribute is successfully read. The caller is responsible for freeing this memory.

**un\_int** Attributes of type **SEC\_INT** and **SEC\_BOOL** store the value of the attribute in this member when the requested attribute is successfully read.

**un\_long**

Attributes of type **SEC\_LONG** store the value of the attribute in this member when the requested attribute is successfully read.

**un\_llong**

Attributes of type **SEC\_LLONG** store the value of the attribute in this member when the requested attribute is successfully read.

**attr\_domain**

The authentication domain containing the attribute. The **getgroupattrs** subroutine is responsible for managing the memory referenced by this pointer.

If **attr\_domain** is specified for an attribute, the get request is sent only to that domain.

If **attr\_domain** is not specified (that is, set to NULL), **getgroupattrs** searches the domains in a predetermined order. The search starts with the local file system and continues with the domains specified in the **/usr/lib/security/methods.cfg** file. This search space can be restricted with the **setauthdb** subroutine so that only the domain specified in the **setauthdb** call is searched.

If **attr\_domain** is not specified, the **getgroupattrs** subroutine sets this field to the name of the domain from which the value is retrieved. If the request for a NULL domain was not satisfied, the request is tried from the local files using the default stanza.

Use the **setuserdb** and **enduserdb** subroutines to open and close the group database. Failure to explicitly open and close the group database can result in loss of memory and performance.

## Parameters

*Group*

Specifies the name of the group for which the attributes are to be read.

*Attributes*

A pointer to an array of 0 or more elements of type **dbattr\_t**. The list of group attributes is defined in the **usersec.h** header file.

*Count* The number of array elements in *Attributes*. A *Count* parameter of 0 can be used to determine if the group exists.

## Security

Files accessed:

Mode	File
rw	/etc/group
rw	/etc/security/group

## Return Values

If *Group* exists, the **getgroupatrrs** subroutine returns 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error. Each element in the *Attributes* array must be examined on a successful call to **getgroupatrrs** to determine if the *Attributes* array entry was successfully retrieved.

## Error Codes

The **getgroupatrrs** subroutine returns an error that indicates that the group does or does not exist. Additional errors can indicate an error querying the information databases for the requested attributes.

<b>EINVAL</b>	The <i>Count</i> parameter is less than zero.
<b>EINVAL</b>	The <i>Attributes</i> parameter is null and the <i>Count</i> parameter is greater than 0.
<b>ENOENT</b>	The specified <i>Group</i> parameter does not exist.

If the **getgroupatrrs** subroutine fails to query an attribute, one or more of the following errors is returned in the **attr\_flag** field of the corresponding *Attributes* element:

<b>EACCES</b>	The user does not have access to the attribute specified in the <i>attr_name</i> field.
<b>EINVAL</b>	The <b>attr_type</b> field in the <i>Attributes</i> entry contains an invalid type.
<b>EINVAL</b>	The <b>attr_un</b> field in the <i>Attributes</i> entry does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors might not be detected.
<b>ENOATTR</b>	The <b>attr_name</b> field in the <i>Attributes</i> entry specifies an attribute that is not defined for this user or group.
<b>ENOMEM</b>	Memory could not be allocated to store the return value.

## Examples

The following sample test program displays the output to a call to **getgroupatrrs**. In this example, the system has a user named foo.

```
attribute name : id
attribute flag : 0
attribute domain : files
attribute value : 204

attribute name : admin
attribute flag : 0
attribute domain : files
attribute value : 0

attribute name : adms
attribute flag : 0
attribute domain : files
attribute value :

attribute name : registry
attribute flag : 0
```

```

attribute domain :
attribute value  : compat

*/
#include <stdio.h>
#include <usersec.h>

#define NATTR 4
#define GROUPNAME "bar"

char * ConvertToComma(char *); /* Convert from SEC_LIST to SEC_CHAR with
                               '\0' replaced with ',' */
main() {

    dbattr_t attributes[NATTR];
    int i;
    int rc;

    memset (&attributes, 0, sizeof(attributes));

    /*
     * Fill in the attributes array with "id", "expires" and
     * "SYSTEM" attributes.
     */

    attributes[0].attr_name = S_ID;
    attributes[0].attr_type = SEC_INT;;

    attributes[1].attr_name = S_ADMIN;
    attributes[1].attr_type = SEC_BOOL;

    attributes[2].attr_name = S_ADMS;
    attributes[2].attr_type = SEC_LIST;

    attributes[3].attr_name = S_REGISTRY;
    attributes[3].attr_type = SEC_CHAR;

    /*
     * Make a call to getuserattrs.
     */
    setuserdb(S_READ);

    rc = getgroupattrs(GROUPNAME, attributes, NATTR);

    enduserdb();

    if (rc) {
        printf("getgroupattrs failed ....\n");
        exit(-1);
    }

    for (i = 0; i < NATTR; i++) {
        printf("attribute name   : %s \n", attributes[i].attr_name);
        printf("attribute flag    : %d \n", attributes[i].attr_flag);

        if (attributes[i].attr_flag) {

            /*
             * No attribute value. Continue.
             */
            printf("\n");
            continue;
        }

        /*
         * We have a value.
         */
        printf("attribute domain : %s \n", attributes[i].attr_domain);
    }
}

```

```

printf("attribute value : ");
switch (attributes[i].attr_type)
{
case SEC_CHAR:
if (attributes[i].attr_char) {
printf("%s\n", attributes[i].attr_char);
free(attributes[i].attr_char);
}
break;
case SEC_LIST:
if (attributes[i].attr_char) {
printf("%s\n", ConvertToComma(
attributes[i].attr_char));
free(attributes[i].attr_char);
}
break;
case SEC_INT:
case SEC_BOOL:
printf("%d\n", attributes[i].attr_int);
break;
default:
break;
}
printf("\n");
}
exit(0);
}

/*
 * ConvertToComme:
 * replaces NULLs in str with commas.
 */
char *
ConvertToComma(char *str)
{
char *s = str;

if (! str || ! *str)
return(s);

for (; *str; str++) {
while(++str);
*str = ',';
}

*(str-1) = 0;
return(s);
}

```

The following output for the call is expected:

```

attribute name : id
attribute flag : 0
attribute domain : files
attribute value : 204

attribute name : admin
attribute flag : 0
attribute domain : files
attribute value : 0

attribute name : adms
attribute flag : 0
attribute domain : files
attribute value :

```

attribute name : registry  
attribute flag : 0  
attribute domain :  
attribute value : compat

## Files

`/etc/group` Contains group IDs.

## Related Information

The “getuserattrs Subroutine” on page 460, `setuserdb` Subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getgroups Subroutine

### Purpose

Gets the supplementary group ID of the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>  
#include <unistd.h>
```

```
int getgroups (NGroups, GIDSet)  
int NGroups;  
gid_t GIDSet [ ];
```

### Description

The **getgroups** subroutine gets the supplementary group ID of the process. The list is stored in the array pointed to by the *GIDSet* parameter. The *NGroups* parameter indicates the number of entries that can be stored in this array. The **getgroups** subroutine never returns more than the number of entries specified by the **NGROUPS\_MAX** constant. (The **NGROUPS\_MAX** constant is defined in the **limits.h** file.) If the value in the *NGroups* parameter is 0, the **getgroups** subroutine returns the number of groups in the supplementary group.

### Parameters

<i>GIDSet</i>	Points to the array in which the supplementary group ID of the user’s process is stored.
<i>NGroups</i>	Indicates the number of entries that can be stored in the array pointed to by the <i>GIDSet</i> parameter.

### Return Values

Upon successful completion, the **getgroups** subroutine returns the number of elements stored into the array pointed to by the *GIDSet* parameter. If the **getgroups** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **getgroups** subroutine is unsuccessful if either of the following error codes is true:

<b>EFAULT</b>	The <i>NGroups</i> and <i>GIDSet</i> parameters specify an array that is partially or completely outside of the allocated address space of the process.
<b>EINVAL</b>	The <i>NGroups</i> parameter is smaller than the number of groups in the supplementary group.

## Related Information

The **getgid** (“getgid, getegid or gegidx Subroutine” on page 370) subroutine, **initgroups** (“initgroups Subroutine” on page 558) subroutine, **setgid** subroutine, **setgroups** subroutine.

The **groups** command, **setgroups** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getgrpaclattr, nextgrpacl, or putgrpaclattr Subroutine

### Purpose

Accesses the group screen information in the SMIT ACL database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>

int getgrpaclattr (Group, Attribute, Value, Type)
char *User;
char *Attribute;
void *Value;
int Type;

char *nextgrpacl(void)

int putgrpaclattr (Group, Attribute, Value, Type)
char *User;
char *Attribute;
void *Value;
int Type;
```

### Description

The **getgrpaclattr** subroutine reads a specified group attribute from the SMIT ACL database. If the database is not already open, this subroutine does an implicit open for reading.

Similarly, the **putgrpaclattr** subroutine writes a specified attribute into the user SMIT ACL database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by the **putgrpaclattr** subroutine must be explicitly committed by calling the **putgrpaclattr** subroutine with a *Type* parameter specifying **SEC\_COMMIT**. Until all the data is committed, only the **getgrpaclattr** subroutine within the process returns written data.

The **nextgrpacl** subroutine returns the next group in a linear search of the group SMIT ACL database. The consistency of consecutive searches depends upon the underlying storage-access mechanism and is not guaranteed by this subroutine.

The **setacldb** and **endacldb** subroutines should be used to open and close the database.



## Parameters

<i>Attribute</i>	Specifies which attribute is read. The following possible attributes are defined in the <b>usersec.h</b> file: <b>S_SCREEN</b> String of SMIT screens. The attribute type is <b>SEC_LIST</b> .
<i>Type</i>	Specifies the type of attribute expected. Valid types are defined in the <b>usersec.h</b> file and include: <b>SEC_LIST</b> The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series must be an empty (zero character count) string.  For the <b>getgrpacltr</b> subroutine, the user should supply a pointer to a defined character pointer variable. For the <b>putgrpacltr</b> subroutine, the user should supply a character pointer.  <b>SEC_COMMIT</b> For the <b>putgrpacltr</b> subroutine, this value specified by itself indicates that changes to the named group are to be committed to permanent storage. The <i>Attribute</i> and <i>Value</i> parameters are ignored. If no group is specified, the changes to all modified groups are committed to permanent storage.  <b>SEC_DELETE</b> The corresponding attribute is deleted from the group SMIT ACL database.  <b>SEC_NEW</b> Updates the group SMIT ACL database file with the new group name when using the <b>putgrpacltr</b> subroutine.
<i>Value</i>	Specifies a buffer, a pointer to a buffer, or a pointer to a pointer depending on the <i>Attribute</i> and <i>Type</i> parameters. See the <i>Type</i> parameter for more details.

## Return Values

If successful, the **getgrpacltr** returns 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

Possible return codes are:

<b>EACCES</b>	Access permission is denied for the data request.
<b>ENOENT</b>	The specified <i>Group</i> parameter does not exist or the attribute is not defined for this group.
<b>ENOATTR</b>	The specified user attribute does not exist for this group.
<b>EINVAL</b>	The <i>Attribute</i> parameter does not contain one of the defined attributes or null.
<b>EINVAL</b>	The <i>Value</i> parameter does not point to a valid buffer or to valid data for this type of attribute.
<b>EPERM</b>	Operation is not permitted.

## Related Information

The **getgrpacltr**, **nextgrpacl**, or **putgrpacltr** (“getgrpacltr, nextgrpacl, or putgrpacltr Subroutine” on page 384) subroutine, **setacldb**, or **endacldb** subroutine.

---

## getgrset Subroutine

### Purpose

Accesses the concurrent group set information in the user database.

## Library

Standard C Library (**libc.a**)

## Syntax

```
char *getgrset (User)
const char * User;
```

## Description

The **getgrset** subroutine returns a pointer to the comma separated list of concurrent group identifiers for the named user.

If the Network Information Service (NIS) is enabled on the system, these subroutines attempt to retrieve the user information from the NIS authentication server.

## Parameters

*User*                    Specifies the user name.

## Return Values

If successful, the **getgrset** subroutine returns a pointer to a list of supplementary groups. This pointer must be freed by the user.

## Error Codes

A **NULL** pointer is returned on error. The value of the **errno** global variable is undefined on error.

## File

*/etc/group*                    Contains basic group attributes.

## Related Information

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine

### Purpose

Manipulates the expiration time of interval timers.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/time.h>
```

```
int getinterval ( TimerID, Value)
timer_t TimerID;
struct itimerstruc_t *Value;
```

```

int incinterval (TimerID, Value, OValue)
timer_t TimerID;
struct itimerstruc_t *Value, *OValue;

int absinterval (TimerID, Value, OValue)
timer_t TimerID;
struct itimerstruc_t *Value, *OValue;

int resabs (TimerID, Resolution, Maximum)
timer_t TimerID;
struct timestruc_t *Resolution, *Maximum;

int resinc (TimerID, Resolution, Maximum)
timer_t TimerID;
struct timestruc_t *Resolution, *Maximum;
#include <unistd.h>

unsigned int alarm ( Seconds)
unsigned int Seconds;

useconds_t ualarm (Value, Interval)
useconds_t Value, Interval;

int setitimer ( Which, Value, OValue)
int Which;
struct itimerval *Value, *OValue;

int getitimer (Which, Value)
int Which;
struct itimerval *Value;

```

## Description

The **getinterval**, **incinterval**, and **absinterval** subroutines manipulate the expiration time of interval timers. These functions use a timer value defined by the **struct itimerstruc\_t** structure, which includes the following fields:

```

struct timestruc_t it_interval; /* timer interval period */
struct timestruc_t it_value; /* timer interval expiration */

```

If the `it_value` field is nonzero, it indicates the time to the next timer expiration. If `it_value` is 0, the per-process timer is disabled. If the `it_interval` member is nonzero, it specifies a value to be used in reloading the `it_value` field when the timer expires. If `it_interval` is 0, the timer is to be disabled after its next expiration (assuming `it_value` is nonzero).

The **getinterval** subroutine returns a value from the **struct itimerstruc\_t** structure to the *Value* parameter. The `it_value` field of this structure represents the amount of time in the current interval before the timer expires, should one exist for the per-process timer specified in the *TimerID* parameter. The `it_interval` field has the value last set by the **incinterval** or **absinterval** subroutine. The fields of the *Value* parameter are subject to the resolution of the timer.

The **incinterval** subroutine sets the value of a per-process timer to a given offset from the current timer setting. The **absinterval** subroutine sets the value of the per-process timer to a given absolute value. If the specified absolute time has already expired, the **absinterval** subroutine will succeed and the expiration notification will be made. Both subroutines update the interval timer period. Time values smaller than the resolution of the specified timer are rounded up to this resolution. Time values larger than the maximum value of the specified timer are rounded down to the maximum value.

The **resinc** and **resabs** subroutines return the resolution and maximum value of the interval timer contained in the *TimerID* parameter. The resolution of the interval timer is contained in the *Resolution* parameter, and the maximum value is contained in the *Maximum* parameter. These values might not be the same as the values returned by the corresponding system timer, the **gettimer** subroutine. In addition, it is likely that the maximum values returned by the **resinc** and **resabs** subroutines will be different.

**Note:** If a nonprivileged user attempts to submit a fine granularity timer (that is, a timer request of less than 10 milliseconds), the timer request is raised to 10 milliseconds.

The **alarm** subroutine causes the system to send the calling thread's process a **SIGALRM** signal after the number of real-time seconds specified by the *Seconds* parameter have elapsed. Since the signal is sent to the process, in a multi-threaded process another thread than the one that called the **alarm** subroutine may receive the **SIGALRM** signal. Processor scheduling delays may prevent the process from handling the signal as soon as it is generated. If the value of the *Seconds* parameter is 0, a pending alarm request, if any, is canceled. Alarm requests are not stacked. Only one **SIGALRM** generation can be scheduled in this manner. If the **SIGALRM** signal has not yet been generated, the call results in rescheduling the time at which the **SIGALRM** signal is generated. If several threads in a process call the **alarm** subroutine, only the last call will be effective.

The **ualarm** subroutine sends a **SIGALRM** signal to the invoking process in a specified number of seconds. The **getitimer** subroutine gets the value of an interval timer. The **setitimer** subroutine sets the value of an interval timer.

## Parameters

<i>TimerID</i>	Specifies the ID of the interval timer.
<i>Value</i>	Points to a <b>struct itimerstruc_t</b> structure.
<i>OValue</i>	Represents the previous time-out period.
<i>Resolution</i>	Resolution of the timer.
<i>Maximum</i>	Indicates the maximum value of the interval timer.
<i>Seconds</i>	Specifies the number of real-time seconds to elapse before the first <b>SIGALRM</b> signal.
<i>Interval</i>	Specifies the number of microseconds between subsequent periodic <b>SIGALRM</b> signals. If a nonprivileged user attempts to submit a fine granularity timer (that is, a timer request of less than 10 milliseconds), the timer request interval is automatically raised to 10 milliseconds.
<i>Which</i>	Identifies the type of timer. Valid values are:

### **ITIMER\_REAL**

Decrements in real time. A **SIGALRM** signal occurs when this timer expires.

### **ITIMER\_VIRTUAL**

Decrements in process virtual time. It runs only during process execution. A **SIGVTALRM** signal occurs when it expires.

### **ITIMER\_PROF**

Decrements in process virtual time and when the system runs on behalf of the process. It is designed for use by interpreters in statistically profiling the execution of interpreted programs. Each time the **ITIMER\_PROF** timer expires, the **SIGPROF** signal occurs. Because this signal may interrupt in-progress system calls, programs using this timer must be prepared to restart interrupted system calls.

## Return Values

If these subroutines are successful, a value of 0 is returned. If an error occurs, a value of -1 is returned and the **errno** global variable is set.

The **alarm** subroutine returns the amount of time (in seconds) remaining before the system is scheduled to generate the **SIGALARM** signal from the previous call to **alarm**. It returns a 0 if there was no previous **alarm** request.

The **ualarm** subroutine returns the number of microseconds previously remaining in the alarm clock.

## Error Codes

If the **getinterval**, **incinterval**, **absinterval**, **resinc**, **resabs**, **setitimer**, **getitimer**, or **setitimer** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one of the following error codes:

<b>EINVAL</b>	Indicates that the <i>TimerID</i> parameter does not correspond to an ID returned by the <b>gettimerid</b> subroutine, or a value structure specified a nanosecond value less than 0 or greater than or equal to one thousand million (1,000,000,000).
<b>EIO</b>	Indicates that an error occurred while accessing the timer device.
<b>EFAULT</b>	Indicates that a parameter address has referenced invalid memory.

The **alarm** subroutine is always successful. No return value is reserved to indicate an error for it.

## Related Information

The **gettimer** (“gettimer, settimer, restimer, stime, or time Subroutine” on page 446) subroutine, **gettimerid** (“gettimerid Subroutine” on page 449) subroutine, **sigaction**, **sigvec**, or **signal** subroutine.

Time data manipulation services in *Operating system and device management*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Signal Management in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* provides more information about signal management in multi-threaded processes.

---

## getiopri Subroutine

### Purpose

Enables the getting of a process I/O priority.

### Syntax

```
short getiopri (ProcessID);  
pid_t ProcessID;
```

### Description

The **getiopri** subroutine returns the I/O scheduling priority of a process. If the target process ID does not match the process ID of the caller, the caller must either be running as root, or have an effective and real user ID that matches the target process.

### Parameters

<i>ProcessID</i>	Specifies the process ID. If this value is -1, the current process scheduling priority is returned.
------------------	---

## Return Values

Upon successful completion, the **getiopri** subroutine returns the I/O scheduling priority of a thread in the process. A returned value of `IOPRIORITY_UNSET` indicates that the I/O priority was not set. Otherwise, a value of `-1` is returned and the **errno** global variable is set to indicate the error.

## Errors

<b>EPERM</b>	The calling process is not root. It does not have the same process ID as the target process, and does not have the same real effective user ID as the target process.
<b>ESRCH</b>	No process can be found corresponding to the specified <i>ProcessID</i> .

## Implementation Specifics

1. Implementation requires an additional field in the **proc** structure.
2. The default setting for process I/O priority is `IOPRIORITY_UNSET`.
3. Once set, process I/O priorities should be inherited across a **fork**. I/O priorities should not be inherited across an **exec**.
4. The **setiopri** system call generates an auditing event using *audit\_svcstart* if auditing is enabled on the system (*audit\_flag* is true).

## Related Information

The “getpri Subroutine” on page 411, **setiopri** subroutine, **setpri** subroutine.

---

## getipnodebyaddr Subroutine

### Purpose

Address-to-nodename translation.

### Library

Standard C Library (**libc.a**)

(**libaixinet**)

### Syntax

```
#include <sys/socket.h>
#include <netdb.h>
struct hostent *getipnodebyaddr(src, len, af, error_num)
const void *src;
size_t len;
int af;
int *error_num;
```

### Description

The **getipnodebyaddr** subroutine has the same arguments as the **gethostbyaddr** subroutine but adds an error number. It is thread-safe.

The **getipnodebyaddr** subroutine is similar in its name query to the **gethostbyaddr** subroutine except in one case. If *af* equals `AF_INET6` and the IPv6 address is an IPv4-mapped IPv6 address or an IPv4-compatible address, then the first 12 bytes are skipped over and the last 4 bytes are used as an IPv4 address with *af* equal to `AF_INET` to lookup the name.

If the **getipnodebyaddr** subroutine is returning success, then the single address that is returned in the **hostent** structure is a copy of the first argument to the function with the same address family and length that was passed as arguments to this function.

All of the information returned by **getipnodebyaddr** is dynamically allocated: the **hostent** structure and the data areas pointed to by the `h_name`, `h_addr_list`, and `h_aliases` members of the **hostent** structure. To return this information to the system the function **freehostent** is called.

## Parameters

<i>src</i>	Specifies a node address. It is a pointer to either a 4-byte (IPv4) or 16-byte (IPv6) binary format address.
<i>af</i>	Specifies the address family which is either AF_INET or AF_INET6.
<i>len</i>	Specifies the length of the node binary format address.
<i>error_num</i>	Returns argument to the caller with the appropriate error code.

## Return Values

The **getipnodebyaddr** subroutine returns a pointer to a **hostent** structure on success.

The **getipnodebyaddr** subroutine returns a null pointer if an error occurs. The *error\_num* parameter is set to indicate the error.

## Error Codes

HOST_NOT_FOUND	The host specified by the <i>name</i> parameter was not found.
TRY_AGAIN	The local server did not receive a response from an authoritative server. Try again later.
NO_RECOVERY	This error code indicates an unrecoverable error.
NO_ADDRESS	The requested <i>name</i> is valid but does not have an Internet address at the name server.

## Related Information

The **freehostent** subroutine and **getipnodebyname** subroutine.

---

## getipnodebyname Subroutine

### Purpose

Nodename-to-address translation.

### Library

Standard C Library (**libc.a**)

(**libaixinet**)

### Syntax

```
#include <libc.a>
#include <netdb.h>
struct hostent *getipnodebyname(name, af, flags, error_num)
```

```
const char *name;
int af;
int flags;
int *error_num;
```

## Description

The commonly used functions **gethostbyname** and **gethostbyname2** are inadequate for many applications. You could not specify the type of addresses desired in **gethostbyname**. In **gethostbyname2**, a global option (RES\_USE\_INET6) is required when IPV6 addresses are used. Also, **gethostbyname2** needed more control over the type of addresses required.

The **getipnodebyname** subroutine gives the caller more control over the types of addresses required and is thread safe. It also does not need a global option like RES\_USE\_INET6.

The name argument can be either a node name or a numeric (either a dotted-decimal IPv4 or colon-separated IPv6) address.

The *flags* parameter values include AI\_DEFAULT, AI\_V4MAPPED, AI\_ALL and AI\_ADDRCONFIG. The special flags value AI\_DEFAULT is designed to handle most applications. Its definition is:

```
#define AI_DEFAULT (AI_V4MAPPED | AI_ADDRCONFIG)
```

When porting simple applications to use IPv6, simply replace the call:

```
hp = gethostbyname(name);
```

with

```
hp = getipnodebyname(name, AF_INET6, AI_DEFAULT, &error_num);
```

To modify the behavior of the **getipnodebyname** subroutine, constant values can be logically-ORed into the *flags* parameter.

A *flags* value of 0 implies a strict interpretation of the *af* parameter. If *af* is AF\_INET then only IPv4 addresses are searched for and returned. If *af* is AF\_INET6 then only IPv6 addresses are searched for and returned.

If the AI\_V4MAPPED flag is specified along with an *af* of AF\_INET6, then the caller accepts IPv4-mapped IPv6 addresses. That is, if a query for IPv6 addresses fails, then a query for IPv4 addresses is made and if any are found, then they are returned as IPv4-mapped IPv6 addresses. The AI\_V4MAPPED flag is only valid with an *af* of AF\_INET6.

If the AI\_ALL flag is used in conjunction the AI\_V4MAPPED flag and *af* is AF\_INET6, then the caller wants all addresses. The addresses returned are IPv6 addresses and/or IPv4-mapped IPv6 addresses. Only if both queries (IPv6 and IPv4) fail does **getipnodebyname** return NULL. Again, the AI\_ALL flag is only valid with an *af* of AF\_INET6.

The AI\_ADDRCONFIG flag is used to specify that a query for IPv6 addresses should only occur if the node has at least one IPv6 source address configured and a query for IPv4 addresses should only occur if the node has at least one IPv4 source address configured. For example, if the node only has IPv4 addresses configured, *af* equals AF\_INET6, and the node name being looked up has both IPv4 and IPv6 addresses, then if only the AI\_ADDRCONFIG flag is specified, **getipnodebyname** will return NULL. If the AI\_V4MAPPED flag is specified with the AI\_ADDRCONFIG flag (AI\_DEFAULT), then any IPv4 addresses found will be returned as IPv4-mapped IPv6 addresses.

There are 4 different situations when the name argument is a literal address string:



1. *name* is a dotted-decimal IPv4 address and *af* is AF\_INET. If the query is successful, then *h\_name* points to a copy of *name*, *h\_addrtype* is the *af* argument, *h\_length* is 4, *h\_aliases* is a NULL pointer, *h\_addr\_list*[0] points to the 4-byte binary address and *h\_addr\_list*[1] is a NULL pointer.
2. *name* is a colon-separated IPv6 address and *af* is AF\_INET6. If the query is successful, then *h\_name* points to a copy of *name*, *h\_addrtype* is the *af* parameter, *h\_length* is 16, *h\_aliases* is a NULL pointer, *h\_addr\_list*[0] points to the 16-byte binary address and *h\_addr\_list*[1] is a NULL pointer.
3. *name* is a dotted-decimal IPv4 address and *af* is AF\_INET6. If the AI\_V4MAPPED flag is specified and the query is successful, then *h\_name* points to an IPv4-mapped IPv6 address string, *h\_addrtype* is the *af* argument, *h\_length* is 16, *h\_aliases* is a NULL pointer, *h\_addr\_list*[0] points to the 16-byte binary address and *h\_addr\_list*[1] is a NULL pointer.
4. *name* is a colon-separated IPv6 address and *af* is AF\_INET. This is an error, **getipnodebyname** returns a NULL pointer and *error\_num* equals HOST\_NOT\_FOUND.

## Parameters

<i>name</i>	Specifies either a node name or a numeric (either a dotted-decimal IPv4 or colon-separated IPv6) address.
<i>af</i>	Specifies the address family which is either AF_INET or AF_INET6.
<i>flags</i>	Controls the types of addresses searched for and the types of addresses returned.
<i>error_num</i>	Returns argument to the caller with the appropriate error code.

## Return Values

The **getipnodebyname** subroutine returns a pointer to a **hostent** structure on success.

The **getipnodebyname** subroutine returns a null pointer if an error occurs. The *error\_num* parameter is set to indicate the error.

## Error Codes

HOST_NOT_FOUND	The host specified by the <i>name</i> parameter was not found.
TRY_AGAIN	The local server did not receive a response from an authoritative server. Try again later.
NO_RECOVERY	The host specified by the <i>name</i> parameter was not found. This error code indicates an unrecoverable error.
NO_ADDRESS	The requested <i>name</i> is valid but does not have an Internet address at the name server.

## Related Information

The **freehostent** subroutine and **getipnodebyaddr** subroutine.

---

## getlogin Subroutine

### Purpose

Gets a user's login name.

### Library

Standard C Library (**libc.a**)

## Syntax

```
include <sys/types.h>
include <unistd.h>
include <limits.h>
char *getlogin (void)
```

## Description

**Attention:** Do not use the **getlogin** subroutine in a multithreaded environment. To access the thread-safe version of this subroutines, see the **getlogin\_r** (“getlogin\_r Subroutine” on page 395) subroutine.

**Attention:** The **getlogin** subroutine returns a pointer to an area that may be overwritten by successive calls.

The **getlogin** subroutine returns a pointer to the login name in the **/etc/utmp** file. You can use the **getlogin** subroutine with the **getpwnam** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422) subroutine to locate the correct password file entry when the same user ID is shared by several login names.

If the **getlogin** subroutine cannot find the login name in the **/etc/utmp** file, it returns the process **LOGNAME** environment variable. If the **getlogin** subroutine is called within a process that is not attached to a terminal, it returns the value of the **LOGNAME** environment variable. If the **LOGNAME** environment variable does not exist, a null pointer is returned.

## Return Values

The return value can point to static data whose content is overwritten by each call. If the login name is not found, the **getlogin** subroutine returns a null pointer.

## Error Codes

If the **getlogin** function is unsuccessful, it returns one or more of the following error codes:

<b>EMFILE</b>	Indicates that the <b>OPEN_MAX</b> file descriptors are currently open in the calling process.
<b>ENFILE</b>	Indicates that the maximum allowable number of files is currently open in the system.
<b>ENXIO</b>	Indicates that the calling process has no controlling terminal.

## Files

**/etc/utmp** Contains a record of users logged into the system.

## Related Information

The **getgrent**, **getgrgid**, **getgrnam**, **putgrent**, **setgrent**, or **endgrent** (“getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371) subroutine, **getlogin\_r** (“getlogin\_r Subroutine” on page 395) subroutine, **getpwent**, **getpwuid**, **setpwent**, or **endpwent** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422) subroutine, **getpwnam** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422) subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getlogin\_r Subroutine

### Purpose

Gets a user's login name.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
int getlogin_r (Name, Length)
char * Name;
size_t Length;
```

### Description

The **getlogin\_r** subroutine gets a user's login name from the **/etc/utmp** file and places it in the *Name* parameter. Only the number of bytes specified by the *Length* parameter (including the ending null value) are placed in the *Name* parameter.

Applications that call the **getlogin\_r** subroutine must allocate memory for the login name before calling the subroutine. The name buffer must be the length of the *Name* parameter plus an ending null value.

If the **getlogin\_r** subroutine cannot find the login name in the **utmp** file or the process is not attached to a terminal, it places the **LOGNAME** environment variable in the name buffer. If the **LOGNAME** environment variable does not exist, the *Name* parameter is set to null and the **getlogin\_r** subroutine returns a -1.

### Parameters

<i>Name</i>	Specifies a buffer for the login name. This buffer should be the length of the <i>Length</i> parameter plus an ending null value.
<i>Length</i>	Specifies the total length in bytes of the <i>Name</i> parameter. No more bytes than the number specified by the <i>Length</i> parameter are placed in the <i>Name</i> parameter, including the ending null value.

### Return Values

If successful, the **getlogin\_r** function returns 0. Otherwise, an error number is returned to indicate the error.

### Error Codes

If the **getlogin\_r** subroutine does not succeed, it returns one of the following error codes:

<b>EINVAL</b>	Indicates that the <i>Name</i> parameter is not valid.
<b>EMFILE</b>	Indicates that the <b>OPEN_MAX</b> file descriptors are currently open in the calling process.
<b>ENFILE</b>	Indicates that the maximum allowable number of files are currently open in the system.
<b>ENXIO</b>	Indicates that the calling process has no controlling terminal.
<b>ERANGE</b>	Indicates that the value of <i>Length</i> is smaller than the length of the string to be returned, including the terminating null character.

### File

<b>/etc/utmp</b>	Contains a record of users logged into the system.
------------------	--

## Related Information

The **getgrent\_r**, **getgrgid\_r**, **getgrnam\_r**, **setgrent\_r**, or **endgrent\_r** (“getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371) subroutine, **getlogin** (“getlogin Subroutine” on page 393) subroutine, **getpwent\_r**, **getpwnam\_r**, **putpwent\_r**, **getpwuid\_r**, **setpwent\_r**, or **endpwent\_r** (“getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine” on page 422) subroutine.

List of Security and Auditing Subroutines, List of Multithread Subroutines, and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getnextprojdb Subroutine

### Purpose

Retrieves the next project from the specified project database.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
getnextprojdb(void *handle, struct project *project, char *comm)
```

### Description

The **getnextprojdb** subroutine retrieves the next project definitions from the project database named through the *handle* parameter. The caller must initialize the project database prior to calling this routine with the **projdballoc** routine. Upon successful completion, the project information is copied to the project structure specified by the caller. In addition, the associated project comment, if present, is copied to the buffer pointed to by the *comm* parameter. The comment buffer is allocated by the caller and must have a length of 1024 bytes.

There is an internal state (that is, the current project) associated with the project database. When the project database is initialized, the current project is the first project in the database. The **getnextprojdb** subroutine returns the current project and advances the current project assignment to the next project in the database so that successive calls read each project entry in the database. When the last project is read, the current project assignment is advanced to the end of the database. Any attempt to read beyond the end of the project database results in a failure.

### Parameters

<i>handle</i>	Pointer to the <b>projdb</b> handle.
<i>project</i>	Pointer to project structure where the retrieved data is stored.
<i>comm</i>	Comment associated with the project in the database.

### Security

No restriction. Any user can call this function.

### Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid arguments, if passed pointer is NULL.
<b>ENOENT</b>	End of the project database.
<b>ENOENT</b>	No projects available.

## Related Information

The “addprojdb Subroutine” on page 32, “chprojattrdb Subroutine” on page 159, “getfirstprojdb Subroutine” on page 368, “getprojdb Subroutine” on page 419, “getprojs Subroutine” on page 420, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, “projdbfree Subroutine” on page 1181, “rmprojdb Subroutine”.

---

## getopt Subroutine

### Purpose

Returns the next flag letter specified on the command line.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int getopt (ArgumentC, ArgumentV, OptionString)
int ArgumentC;
char *const ArgumentV [ ];
const char *OptionString;
```

```
extern int optind;
extern int optopt;
extern int opterr;
extern char *optarg;
```

### Description

The *optind* parameter indexes the next element of the *ArgumentV* parameter to be processed. It is initialized to 1 and the **getopt** subroutine updates it after calling each element of the *ArgumentV* parameter.

The **getopt** subroutine returns the next flag letter in the *ArgumentV* parameter list that matches a letter in the *OptionString* parameter. If the flag takes an argument, the **getopt** subroutine sets the *optarg* parameter to point to the argument as follows:

- If the flag was the last letter in the string pointed to by an element of the *ArgumentV* parameter, the *optarg* parameter contains the next element of the *ArgumentV* parameter and the *optind* parameter is incremented by 2. If the resulting value of the *optind* parameter is not less than the *ArgumentC* parameter, this indicates a missing flag argument, and the **getopt** subroutine returns an error message.
- Otherwise, the *optarg* parameter points to the string following the flag letter in that element of the *ArgumentV* parameter and the *optind* parameter is incremented by 1.

### Parameters

*ArgumentC* Specifies the number of parameters passed to the routine.

<i>ArgumentV</i>	Specifies the list of parameters passed to the routine.
<i>OptionString</i>	Specifies a string of recognized flag letters. If a letter is followed by a : (colon), the flag is expected to take a parameter that may or may not be separated from it by white space.
<i>optind</i>	Specifies the next element of the <i>ArgumentV</i> array to be processed.
<i>optopt</i>	Specifies any erroneous character in the <i>OptionString</i> parameter.
<i>opterr</i>	Indicates that an error has occurred when set to a value other than 0.
<i>optarg</i>	Points to the next option flag argument.

## Return Values

The **getopt** subroutine returns the next flag letter specified on the command line. A value of -1 is returned when all command line flags have been parsed. When the value of the *ArgumentV* [*optind*] parameter is null, *\*ArgumentV* [*optind*] is not the - (minus) character, or *ArgumentV* [*optind*] points to the "-" (minus) string, the **getopt** subroutine returns a value of -1 without changing the value. If *ArgumentV* [*optind*] points to the "- -" (double minus) string, the **getopt** subroutine returns a value of -1 after incrementing the value of the *optind* parameter.

## Error Codes

If the **getopt** subroutine encounters an option character that is not specified by the *OptionString* parameter, a ? (question mark) character is returned. If it detects a missing option argument and the first character of *OptionString* is a : (colon), then a : (colon) character is returned. If this subroutine detects a missing option argument and the first character of *OptionString* is not a colon, it returns a ? (question mark). In either case, the **getopt** subroutine sets the *optopt* parameter to the option character that caused the error. If the application has not set the *opterr* parameter to 0 and the first character of *OptionString* is not a : (colon), the **getopt** subroutine also prints a diagnostic message to standard error.

## Examples

The following code fragment processes the flags for a command that can take the mutually exclusive flags **a** and **b**, and the flags **f** and **o**, both of which require parameters.

```
#include <unistd.h> /*Needed for access subroutine constants*/
main (argc, argv)
int argc;
char **argv;
{
    int c;
    extern int optind;
    extern char *optarg;
    .
    .
    .
    while ((c = getopt(argc, argv, "abf:o:")) != EOF)
    {
        switch (c)
        {
            case 'a':
                if (bflg)
                    errflg++;
                else
                    aflg++;
                break;
            case 'b':
                if (aflg)
                    errflg++;
                else
                    bflg++;
                break;
```

```

        case 'f':
            ifile = optarg;
            break;
        case 'o':
            ofile = optarg;
            break;
        case '?':
            errflg++;
    } /* case */
    if (errflg)
    {
        fprintf(stderr, "usage: . . . ");
        exit(2);
    }
} /* while */
for ( ; optind < argc; optind++)
{
    if (access(argv[optind], R_OK))
    {
        .
        .
        .
    }
} /* for */
} /* main */

```

## Related Information

The **getopt** command.

List of Executable Program Creation Subroutines, Subroutines Overview, and List of Multithread Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpagesize Subroutine

### Purpose

Gets the system page size.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int getpagesize( )
```

### Description

The **getpagesize** subroutine returns the number of bytes in a page. Page granularity is the granularity for many of the memory management calls.

The page size is determined by the system and may not be the same as the underlying hardware page size.

## Related Information

The **brk** or **sbrk** (“brk or sbrk Subroutine” on page 122) subroutine.

The **pagesize** command.

Program Address Space Overview and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpinfo Subroutine

### Purpose

Retrieves a Process Authentication Group (PAG) flags for a given PAG type.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <pag.h>
```

```
int getpinfo ( name, infop, infosz )
char * name;
struct paginfo * infop;
int infosz;
```

### Description

The **getpinfo** subroutine retrieves the PAG flags for a given PAG name. For this function to succeed, the PAG name must be registered with the operating system before this subroutine is called. The *infop* parameter must be a valid, referenced PAG info structure of the size specified by *infosz*.

### Parameters

<i>name</i>	A 1-character to 4-character, NULL-terminated name for the PAG type. Typical values include afs, dfs, pki, and krb5.
<i>infop</i>	Points to a <b>paginfo struct</b> where the operating system returns the PAG flags.
<i>infosz</i>	Indicates the size of the PAG info structure.

### Return Values

A value of 0 is returned upon successful completion. If the **getpinfo** subroutine fails a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **getpinfo** subroutine fails if the following condition is true:

<b>EINVAL</b>	The named PAG type does not exist as part of the table.
---------------	---

Other errors might be set by subroutines invoked by the **getpinfo** subroutine.

## Related Information

**\_\_pag\_getid** System Call, **\_\_pag\_getname** System Call, **\_\_pag\_getvalue** System Call, **\_\_pag\_setname** System Call, **\_\_pag\_setvalue** System Call, **kcred\_genpagvalue** Kernel Service, **kcred\_getpagid** Kernel Service, and **kcred\_getpagname** Kernel Service.



## getpagvalue or getpagvalue64 Subroutine

### Purpose

Returns the Process Authentication Group (PAG) value for a given PAG type.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <pag.h>
```

```
int getpagvalue ( name )  
char * name;
```

```
uint64_t getpagvalue64( name );  
char * name;
```

### Description

The **getpagvalue** and **getpagvalue64** subroutines retrieve the PAG value for a given PAG name. For these functions to succeed, the PAG name must be registered with the operating system before these subroutines are called.

### Parameters

*name*            A 1-character to 4-character, NULL-terminated name for the PAG type. Typical values include afs, dfs, pki, and krb5.

### Return Values

The **getpagvalue** and **getpagvalue64** subroutines return a PAG value upon successful completion. Upon a failure, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **getpagvalue** and **getpagvalue64** subroutines fail if the following condition is true:

**EINVAL**                    The named PAG type does not exist as part of the table.

Other errors might be set by subroutines invoked by the **getpagvalue** and **getpagvalue64** subroutines.

### Related Information

**\_\_pag\_getid** System Call, **\_\_pag\_getname** System Call, **\_\_pag\_getvalue** System Call, **\_\_pag\_setname** System Call, **\_\_pag\_setvalue** System Call, **kcred\_genpagvalue** Kernel Service, **kcred\_getpagid** Kernel Service, and **kcred\_getpagname** Kernel Service.

List of Security and Auditing Subroutines in *AIX 5L Version 5.3 General Programming Concepts*.

---

## getpass Subroutine

### Purpose

Reads a password.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
char *getpass ( Prompt)  
char *Prompt;
```

### Description

**Attention:** The characters are returned in a static data area. Subsequent calls to this subroutine overwrite the static data area.

The **getpass** subroutine does the following:

- Opens the controlling terminal of the current process.
- Writes the characters specified by the *Prompt* parameter to that device.
- Reads from that device the number of characters up to the value of the **PASS\_MAX** constant until a new-line or end-of-file (EOF) character is detected.
- Restores the terminal state and closes the controlling terminal.

During the read operation, character echoing is disabled.

The **getpass** subroutine is not safe in a multithreaded environment. To use the **getpass** subroutine in a threaded application, the application must keep the integrity of each thread.

### Parameters

*Prompt* Specifies a prompt to display on the terminal.

### Return Values

If this subroutine is successful, it returns a pointer to the string. If an error occurs, the subroutine returns a null pointer and sets the **errno** global variable to indicate the error.

### Error Codes

If the **getpass** subroutine is unsuccessful, it returns one or more of the following error codes:

- |              |  |
|--------------|--|
| <b>EINTR</b> | Indicates that an interrupt occurred while the <b>getpass</b> subroutine was reading the terminal device. If a <b>SIGINT</b> or <b>SIGQUIT</b> signal is received, the <b>getpass</b> subroutine terminates input and sends the signal to the calling process. |
| <b>ENXIO</b> | Indicates that the process does not have a controlling terminal.   |

**Note:** Any subroutines called by the **getpass** subroutine may set other error codes.

## Related Information

The **getuserpw** (“getuserpw, putuserpw, or putuserpwhist Subroutine” on page 468) subroutine, **newpass** (“newpass Subroutine” on page 901) subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpcred Subroutine

### Purpose

Reads the current process credentials.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
char **getpcred ( Which)  
int Which;
```

### Description

The **getpcred** subroutine reads the specified process security credentials and returns a pointer to a NULL terminated array of pointers in allocated memory. Each pointer in the array points to a string containing an attribute/value pair in allocated memory. It's the responsibility of the caller to free each individual string as well as the array of pointers.

## Parameters

*Which* Specifies which credentials are read. This parameter is a bit mask and can contain one or more of the following values, as defined in the **usersec.h** file:

**CRED\_RUID**

Real user name

**CRED\_LUID**

Login user name

**CRED\_RGID**

Real group name

**CRED\_GROUPS**

Supplementary group ID

**CRED\_AUDIT**

Audit class of the current process

**Note:** A process must have root user authority to retrieve this credential. Otherwise, the **getpcred** subroutine returns a null pointer and the **errno** global variable is set to **EPERM**.

**CRED\_RLIMITS**

BSD resource limits

**Note:** Use the **getrlimit** (“getrlimit, getrlimit64, setrlimit, setrlimit64, or vlimit Subroutine” on page 424) subroutine to control resource consumption.

**CRED\_UMASK**

The umask.

If the *Which* parameter is null, all credentials are returned.

## Return Values

When successful, the **getpcred** subroutine returns a pointer to a NULL terminated array of string pointers containing the requested values. If the **getpcred** subroutine is unsuccessful, a NULL pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **getpcred** subroutine fails if either of the following are true:

**EINVAL** The *Which* parameter contains invalid credentials requests.

**EPERM** The process does not have the proper authority to retrieve the requested credentials.

Other errors can also be set by any subroutines invoked by the **getpcred** subroutine.

## Related Information

The **ckuseracct** (“ckuseracct Subroutine” on page 164) subroutine, **ckuserID** (“ckuserID Subroutine” on page 166) subroutine, **getpenv** (“getpenv Subroutine” on page 405) subroutine, **setpenv** subroutine, **setpcred** subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpeereid Subroutine

### Purpose

Gets the effective user ID and effective group ID of a peer on a connected UNIX domain socket.

### Syntax

```
#include <sys/types.h>
int getpeereid (int socket, uid_t *euid, gid_t *egid)
```

### Description

The **getpeereid** subroutine returns the effective user and group IDs of the peer connected to a stream socket in the UNIX domain. The effective user and group IDs are saved in the socket, to be returned, when the peer calls **connect** or **listen**.

### Parameters

<i>socket</i>	Specifies the descriptor number of a connected socket.
<i>euid</i>	The effective user ID of the peer socket.
<i>egid</i>	The effective group ID of the peer socket.

### Return Values

When the **getpeereid** subroutine successfully completes, a value of 0 is returned and the *euid* and *egid* parameters hold the effective user ID and group ID, respectively.

If the **getpeereid** subroutine is unsuccessful, the system handler returns a value of -1 to the calling program and sets the **errno** global variable to an error code that indicates the specific error.

### Error Codes

The **getpeereid** subroutine is unsuccessful if any of the following errors occurs:

<b>EBADF</b>	The <i>socket</i> parameter is not valid.
<b>ENOTSOCK</b>	The <i>socket</i> parameter refers to a file, not a socket.
<b>ENOTCONN</b>	The socket is not connected.
<b>ENOBUFS</b>	Insufficient resources were available in the system to complete the call.
<b>EFAULT</b>	The <i>address</i> parameter is not in a writable part of the user address space.

**Note:** The **getpeerid** technology used to support this function in AIX was originally published by D. J. Bernstein, Associate Professor, Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago. In addition, the specific **getpeerid** syntax reflected originated with William Erik Baxter. All the aforementioned are used by AIX with permission.

---

## getpenv Subroutine

### Purpose

Reads the current process environment.

### Library

Security Library (**libc.a**)

## Syntax

```
#include <usersec.h>
```

```
char **getpenv ( Which)  
int Which;
```

## Description

The **getpenv** subroutine reads the specified environment variables and returns them in a character buffer.

## Parameters

*Which* Specifies which environment variables are to be returned. This parameter is a bit mask and may contain one or more of the following values, as defined in the **usersec.h** file:

### **PENV\_USR**

The normal user-state environment. Typically, the shell variables are contained here.

### **PENV\_SYS**

The system-state environment. This data is located in system space and protected from unauthorized access.

All variables are returned by setting the *Which* parameter to logically OR the **PENV\_USER** and **PENV\_SYSTEM** values.

The variables are returned in a null-terminated array of character pointers in the form `var=val`. The user-state environment variables are prefaced by the string **USRENVIRON:**, and the system-state variables are prefaced with **SYSENVIRON:**. If a user-state environment is requested, the current directory is always returned in the **PWD** variable. If this variable is not present in the existing environment, the **getpenv** subroutine adds it to the returned string.

## Return Values

Upon successful return, the **getpenv** subroutine returns the environment values. If the **getpenv** subroutine fails, a null value is returned and the **errno** global variable is set to indicate the error.

**Note:** This subroutine can partially succeed, returning only the values that the process permits it to read.

## Error Codes

The **getpenv** subroutine fails if one or more of the following are true:

**EINVAL** The *Which* parameter contains values other than **PENV\_USR** or **PENV\_SYS**.

Other errors can also be set by subroutines invoked by the **getpenv** subroutine.

## Related Information

The **ckuseracct** (“ckuseracct Subroutine” on page 164) subroutine, **ckuserID** (“ckuserID Subroutine” on page 166) subroutine, **getpcred** (“getpcred Subroutine” on page 403) subroutine, **setpenv** subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpgid Subroutine

### Purpose

Returns the process group ID of the calling process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
pid_t getpgid (Pid)
(pid_ Pid)
```

### Description

The **getpgid** subroutine returns the process group ID of the process whose process ID is equal to that specified by the *Pid* parameter. If the value of the *Pid* parameter is equal to **(pid\_t)0**, the **getpgid** subroutine returns the process group ID of the calling process.

### Parameter

*Pid*                    The process ID of the process to return the process group ID for.

### Return Values

id                    The process group ID of the requested process  
-1                    Not successful and **errno** set to one of the following.

### Error Code

<b>ESRCH</b>	There is no process with a process ID equal to <i>Pid</i> .
<b>EPERM</b>	The process whose process ID is equal to <i>Pid</i> is not in the same session as the calling process.
<b>EINVAL</b>	The value of the <i>Pid</i> argument is invalid.

### Related Information

The **exec** (“**exec**: **execl**, **execle**, **execlp**, **execv**, **execve**, **execvp**, or **exec** Subroutine” on page 236) subroutine, **fork** (“**fork**, **f\_fork**, or **vfork** Subroutine” on page 291) subroutine, **getpid** (“**getpid**, **getpgrp**, or **getppid** Subroutine”) subroutine, **getsid** (“**getsid** Subroutine” on page 436) subroutine, **setpgid** subroutine, **setsid** subroutine.

---

## getpid, getpgrp, or getppid Subroutine

### Purpose

Returns the process ID, process group ID, and parent process ID.

## Syntax

```
#include <unistd.h>
pid_t getpid (void)
pid_t getpgrp (void)
pid_t getppid (void)
```

## Description

The **getpid** subroutine returns the process ID of the calling process.

The **getpgrp** subroutine returns the process group ID of the calling process.

The **getppid** subroutine returns the process ID of the calling process' parent process.

## Related Information

The **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutines, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **setpgid** subroutine, **setpgrp** subroutine, **sigaction**, **sigvec**, or **signal** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getportattr or putportattr Subroutine

### Purpose

Accesses the port information in the port database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>

int getportattr (Port, Attribute, Value, Type)
char * Port;
char * Attribute;
void * Value;
int Type;

int putportattr (Port, Attribute, Value, Type)
char *Port;
char *Attribute;
void *Value;
int Type;
```

### Description

The **getportattr** or **putportattr** subroutine accesses port information. The **getportattr** subroutine reads a specified attribute from the port database. If the database is not already open, the **getportattr** subroutine implicitly opens the database for reading. The **putportattr** subroutine writes a specified attribute into the port database. If the database is not already open, the **putportattr** subroutine implicitly opens the database for reading and writing. The data changed by the **putportattr** subroutine must be explicitly committed by calling the **putportattr** subroutine with a *Type* parameter equal to the **SEC\_COMMIT** value. Until all the data is committed, only these subroutines within the process return the written data.



Values returned by these subroutines are in dynamically allocated buffers. You do not need to move the values prior to the next call.

Use the **setuserdb** or **enduserdb** subroutine to open and close the port database.

## Parameters

<i>Port</i>	Specifies the name of the port for which an attribute is read.
<i>Attribute</i>	Specifies the name of the attribute read. This attribute can be one of the following values defined in the <b>usersec.h</b> file:  <b>S_HERALD</b> Defines the initial message printed when the <b>getty</b> or <b>login</b> command prompts for a login name. This value is of the type <b>SEC_CHAR</b> .  <b>S_SAKENABLED</b> Indicates whether or not trusted path processing is allowed on this port. This value is of the type <b>SEC_BOOL</b> .  <b>S_SYNONYM</b> Defines the set of ports that are <b>synonym</b> attributes for the given port. This value is of the type <b>SEC_LIST</b> .  <b>S_LOGTIMES</b> Defines when the user can access the port. This value is of the type <b>SEC_LIST</b> .  <b>S_LOGDISABLE</b> Defines the number of unsuccessful login attempts that result in the system locking the port. This value is of the type <b>SEC_INT</b> .  <b>S_LOGINTERVAL</b> Defines the time interval in seconds within which <b>S_LOGDISABLE</b> number of unsuccessful login attempts must occur before the system locks the port. This value is of the type <b>SEC_INT</b> .  <b>S_LOGREENABLE</b> Defines the time interval in minutes after which a system-locked port is unlocked. This value is of the type <b>SEC_INT</b> .  <b>S_LOGDELAY</b> Defines the delay factor in seconds between unsuccessful login attempts. This value is of the type <b>SEC_INT</b> .  <b>S_LOCKTIME</b> Defines the time in seconds since the epoch (zero time, January 1, 1970) that the port was locked. This value is of the type <b>SEC_INT</b> .  <b>S_ULOGTIMES</b> Lists the times in seconds since the epoch (midnight, January 1, 1970) when unsuccessful login attempts occurred. This value is of the type <b>SEC_LIST</b> .  <b>S_USERNAMEECHO</b> Indicates whether user name input echo and user name masking is enabled for the port. This value is of the type <b>SEC_BOOL</b> .  <b>S_PWD PROMPT</b> Defines the password prompt message printed when requesting password input. This value is of the type <b>SEC_CHAR</b> .
<i>Value</i>	Specifies the address of a buffer in which the attribute is stored with <b>putportattr</b> or is to be read <b>getportattr</b> .

<i>Type</i>	Specifies the type of attribute expected. The following types are valid and defined in the <b>usersec.h</b> file:
<b>SEC_INT</b>	Indicates the format of the attribute is an integer. The buffer returned by the <b>getportattr</b> subroutine and the buffer supplied by the <b>putportattr</b> subroutine are defined to contain an integer.
<b>SEC_CHAR</b>	Indicates the format of the attribute is a null-terminated character string.
<b>SEC_LIST</b>	Indicates the format of the attribute is a list of null-terminated character strings. The list itself is null terminated.
<b>SEC_BOOL</b>	An integer with a value of either 0 or 1, or a pointer to a character pointing to one of the following strings: <ul style="list-style-type: none"> <li>• True</li> <li>• Yes</li> <li>• Always</li> <li>• False</li> <li>• No</li> <li>• Never</li> </ul>
<b>SEC_COMMIT</b>	Indicates that changes to the specified port are committed to permanent storage if specified alone for the <b>putportattr</b> subroutine. The <i>Attribute</i> and <i>Value</i> parameters are ignored. If no port is specified, changes to all modified ports are committed.
<b>SEC_DELETE</b>	Deletes the corresponding attribute from the database.
<b>SEC_NEW</b>	Updates all of the port database files with the new port name when using the <b>putportattr</b> subroutine.

## Security

Access Control: The calling process must have access to the port information in the port database.

File Accessed:

<b>rw</b>	/etc/security/login.cfg
<b>rw</b>	/etc/security/portlog

## Return Values

The **getportattr** and **putportattr** subroutines return a value of 0 if completed successfully. Otherwise, a value of -1 is returned and the **errno** global value is set to indicate the error.

## Error Codes

These subroutines are unsuccessful if the following values are true:

<b>EACCES</b>	Indicates that access permission is denied for the data requested.
<b>ENOENT</b>	Indicates that the <i>Port</i> parameter does not exist or the attribute is not defined for the specified port.
<b>ENOATTR</b>	Indicates that the specified port attribute does not exist for the specified port.

<b>EINVAL</b>	Indicates that the <i>Attribute</i> parameter does not contain one of the defined attributes or is a null value.
<b>EINVAL</b>	Indicates that the <i>Value</i> parameter does not point to a valid buffer or to valid data for this type of attribute.
<b>EPERM</b>	Operation is not permitted.

## Related Information

The **setuserdb** or **enduserdb** subroutine.

List of Security and Auditing Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpri Subroutine

### Purpose

Returns the scheduling priority of a process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int getpri ( ProcessID)
pid_t ProcessID;
```

### Description

The **getpri** subroutine returns the scheduling priority of a process.

### Parameters

*ProcessID* Specifies the process ID. If this value is 0, the current process scheduling priority is returned.

### Return Values

Upon successful completion, the **getpri** subroutine returns the scheduling priority of a thread in the process. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **getpri** subroutine is unsuccessful if one of the following is true:

<b>EPERM</b>	A process was located, but its effective and real user ID did not match that of the process running the <b>getpri</b> subroutine, and the calling process did not have root user authority.
<b>ESRCH</b>	No process can be found corresponding to that specified by the <i>ProcessID</i> parameter.

## Related Information

The **setpri** subroutine.

Performance-Related Subroutines in *Performance management*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getpriority, setpriority, or nice Subroutine

### Purpose

Gets or sets the nice value.

### Libraries

**getpriority**, **setpriority**: Standard C Library (**libc.a**)

**nice**: Standard C Library (**libc.a**)

Berkeley Compatibility Library (**libbsd.a**)

### Syntax

```
#include <sys/resource.h>
```

```
int getpriority( Which, Who)
int Which;
int Who;
```

```
int setpriority(Which, Who, Priority)
int Which;
int Who;
int Priority;
```

```
#include <unistd.h>
```

```
int nice( Increment)
int Increment;
```

### Description

The nice value of the process, process group, or user, as indicated by the *Which* and *Who* parameters is obtained with the **getpriority** subroutine and set with the **setpriority** subroutine.

The **getpriority** subroutine returns the highest priority nice value (lowest numerical value) pertaining to any of the specified processes. The **setpriority** subroutine sets the nice values of all of the specified processes to the specified value. If the specified value is less than -20, a value of -20 is used; if it is greater than 20, a value of 20 is used. Only processes that have root user authority can lower nice values.

The **nice** subroutine increments the nice value by the value of the *Increment* parameter.

**Note:** Nice values are only used for the scheduling policy **SCHED\_OTHER**, where they are combined with a calculation of recent cpu usage to determine the priority value.

To provide upward compatibility with older programs, the **nice** interface, originally found in AT&T System V, is supported.

**Note:** Process priorities in AT&T System V are defined in the range of 0 to 39, rather than -20 to 20 as in BSD, and the **nice** library routine is supported by both. Accordingly, two versions of the **nice** are supported by AIX Version 3. The default version behaves like the AT&T System V version, with the *Increment* parameter treated as the modifier of a value in the range of 0 to 39 (0 corresponds to -20, 39 corresponds to 9, and priority 20 is not reachable with this interface).

If the behavior of the BSD version is desired, compile with the Berkeley Compatibility Library (**libbsd.a**). The *Increment* parameter is treated as the modifier of a value in the range -20 to 20.

## Parameters

<i>Which</i>	Specifies one of <b>PRIO_PROCESS</b> , <b>PRIO_PGRP</b> , or <b>PRIO_USER</b> .
<i>Who</i>	Interpreted relative to the <i>Which</i> parameter (a process identifier, process group identifier, and a user ID, respectively). A zero value for the <i>Who</i> parameter denotes the current process, process group, or user.
<i>Priority</i>	Specifies a value in the range -20 to 20. Negative nice values cause more favorable scheduling.
<i>Increment</i>	Specifies a value that is added to the current process nice value. Negative values can be specified, although values exceeding either the high or low limit are truncated.

## Return Values

On successful completion, the **getpriority** subroutine returns an integer in the range -20 to 20. A return value of -1 can also indicate an error, and in this case the **errno** global variable is set.

On successful completion, the **setpriority** subroutine returns 0. Otherwise, -1 is returned and the global variable **errno** is set to indicate the error.

On successful completion, the **nice** subroutine returns the new nice value minus {NZERO}. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

**Note:** A value of -1 can also be returned. In that case, the calling process should also check the **errno** global variable.

## Error Codes

The **getpriority** and **setpriority** subroutines are unsuccessful if one of the following is true:

<b>ESRCH</b>	No process was located using the <i>Which</i> and <i>Who</i> parameter values specified.
<b>EINVAL</b>	The <i>Which</i> parameter was not recognized.

In addition to the errors indicated above, the **setpriority** subroutine is unsuccessful if one of the following is true:

<b>EPERM</b>	A process was located, but neither the effective nor real user ID of the caller of the process executing the <b>setpriority</b> subroutine has root user authority.
<b>EACCES</b>	The call to <b>setpriority</b> would have changed the priority of a process to a value lower than its current value, and the effective user ID of the process executing the call did not have root user authority.

The **nice** subroutine is unsuccessful if the following is true:

<b>EPERM</b>	The <i>Increment</i> parameter is negative or greater than $2 * \{\text{NZERO}\}$ and the calling process does not have appropriate privileges.
--------------	---

## Related Information

The **exec** (“exec: execl, execlp, execlp, execv, execve, execvp, or exact Subroutine” on page 236) subroutines.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getproclist, getlparlist, or getarmlist Subroutine

### Purpose

Retrieve the transaction records from the advanced accounting data file.

### Library

The **libaacct.a** library.

### Syntax

```
#include <sys/aacct.h>
int getproclist(filename, begin_time, end_time, p_list);
int getlparlist(filename, begin_time, end_time, l_list);
int getarmlist(filename, begin_time, end_time, t_list);
char *filename;
long long begin_time;
long long end_time;
struct aaacct_tran **p_list, **l_list, **t_list
```

### Description

The **getproclist**, **getlparlist**, and **getarmlist** subroutines parse the specified advanced accounting data file and retrieve the process, LPAR, and ARM transaction records, respectively. The retrieved transaction records are returned in the form of a linked list of type **struct aaacct\_tran\_rec**.

These APIs can be called multiple times with different accounting data file names in order to generate a consolidated list of transaction records from multiple data files. They append the new file data to the end of the linked list pointed to by the *p\_list*, *l\_list*, and *t\_list* arguments. They also internally sort the transaction records based on the time of transaction, which gives users a time-sorted list of transaction records from these routines.

The **getproclist**, **getlparlist**, and **getarmlist** subroutines can also be used to retrieve the intended transaction records for a particular interval of time by passing the begin and end times of the interval as arguments to these routines. If these interval arguments are specified as -1, transaction records for all the intervals are retrieved.

### Parameters

<i>begin_time</i>	Specifies the start timestamp for collecting records in a particular intervals. The input is in seconds since EPOCH. Specifying -1 retrieves all the records.
<i>end_time</i>	Specifies the end timestamp for collecting records in a particular intervals. The input is in seconds since EPOCH. Specifying -1 retrieves all the records.
<i>filename</i>	Name of the advanced accounting data file.
<i>l_list</i>	Pointers to the linked list of <b>aaacct_tran_rec</b> structures, which hold the retrieved LPAR records.
<i>p_list</i>	Pointers to the linked list of <b>aaacct_tran_rec</b> structures, which hold the retrieved process records.
<i>t_list</i>	Pointers to the linked list of <b>aaacct_tran_rec</b> structures, which hold the retrieved ARM records.

## Security

No restrictions. Any user can call this function.

## Return Values

0	The call to the subroutine was successful.
-1	The call to the subroutine failed.

## Error Codes

<b>EINVAL</b>	The passed pointer is NULL.
<b>ENOENT</b>	Specified data file does not exist.
<b>EPERM</b>	Permission denied. Unable to read the data file.
<b>ENOMEM</b>	Insufficient memory.

## Related Information

The “agg\_proc\_stat, agg\_lpar\_stat, agg\_arm\_stat, or free\_agg\_list Subroutine” on page 36, “buildproclist Subroutine” on page 125, “buildtranlist or freetranlist Subroutine” on page 126.

Understanding the Advanced Accounting Subsystem.

---

## getprocs Subroutine

### Purpose

Gets process table entries.

### Library

Standard C library (**libc.a**)

## Syntax

```
#include <procsinfo.h>
#include <sys/types.h>
```

```
int
getprocs ( ProcessBuffer, ProcessSize, FileBuffer, FileSize, IndexPointer, Count)
struct procsinfo *ProcessBuffer;
or struct procsinfo64 *ProcessBuffer;
int ProcessSize;
struct fdsinfo *FileBuffer;
int FileSize;
pid_t *IndexPointer;
int Count;
```

```
int
getprocs64 ( ProcessBuffer, ProcessSize, FileBuffer, FileSize, IndexPointer, Count)
struct procsinfo64 *ProcessBuffer;
int ProcessSize;
struct fdsinfo64 *FileBuffer;
int FileSize;
pid_t *IndexPointer;
int Count;
```

## Description

The **getprocs** subroutine returns information about processes, including process table information defined by the **procsinfo** structure, and information about the per-process file descriptors defined by the **fdsinfo** structure.

The **getprocs** subroutine retrieves up to *Count* process table entries, starting with the process table entry corresponding to the process identifier indicated by *IndexPointer*, and places them in the array of **procsinfo** structures indicated by the *ProcessBuffer* parameter. File descriptor information corresponding to the retrieved processes are stored in the array of **fdsinfo** structures indicated by the *FileBuffer* parameter.

On return, the process identifier referenced by *IndexPointer* is updated to indicate the next process table entry to be retrieved. The **getprocs** subroutine returns the number of process table entries retrieved.

The **getprocs** subroutine is normally called repeatedly in a loop, starting with a process identifier of zero, and looping until the return value is less than *Count*, indicating that there are no more entries to retrieve.

**Note:** The process table may change while the **getprocs** subroutine is accessing it. Returned entries will always be consistent, but since processes can be created or destroyed while the **getprocs** subroutine is running, there is no guarantee that retrieved entries will still exist, or that all existing processes have been retrieved.

When used in 32-bit mode, limits larger than can be represented in 32 bits are truncated to RLIM\_INFINITY. Large **rusage** and other values are truncated to INT\_MAX. Alternatively, the **struct procsinfo64** and *sizeof(struct procsinfo64)* can be used by 32-bit **getprocs** to return full 64-bit process information. Note that the **procsinfo** structure not only increases certain **procsinfo** fields from 32 to 64 bits, but that it contains additional information not present in **procsinfo**. The **struct procsinfo64** contains the same data as **struct procsinfo** when compiled in a 64-bit program.



In AIX 5.1 and later, 64-bit applications are required to use **getprocs64()** and **procentry64**. Note that **struct procentry64** contains the same information as **struct procsinfo64**, with the addition of support for the 64-bit `time_t` and `dev_t`, and the 256-bit `sigset_t`. The **procentry64** structure also contains a new version of **struct ucred** (**struct ucred\_ext**) and a new, expanded **struct rusage** (**struct trusage64**) as described in `<sys/cred.h>` and `<sys/resource.h>` respectively. Application developers are also encouraged to use **getprocs64()** in 32-bit applications to obtain 64-bit process information as this interface provides the new, larger types. The **getprocs()** interface will still be supported for 32-bit applications using **struct procsinfo** or **struct procsinfo64** but will not be available to 64-bit applications.

## Parameters

### *ProcessBuffer*

Specifies the starting address of an array of **procsinfo**, **procsinfo64**, or **procentry64** structures to be filled in with process table entries. If a value of **NULL** is passed for this parameter, the **getprocs** subroutine scans the process table and sets return values as normal, but no process entries are retrieved.

**Note:** The *ProcessBuffer* parameter of **getprocs** subroutine contains two struct `rusage` fields named **pi\_ru** and **pi\_cru**. Each of these fields contains two struct `timeval` fields named **ru\_utime** and **ru\_stime**. The **tv\_usec** field in both of the struct `timeval` contain nanoseconds instead of microseconds. These values come from the struct user fields named **U\_ru** and **U\_cru**.

### *ProcessSize*

Specifies the size of a single **procsinfo**, **procsinfo64**, or **procentry64** structure.

### *FileBuffer*

Specifies the starting address of an array of **fdsinfo**, or **fdsinfo64** structures to be filled in with per-process file descriptor information. If a value of **NULL** is passed for this parameter, the **getprocs** subroutine scans the process table and sets return values as normal, but no file descriptor entries are retrieved.

### *FileSize*

Specifies the size of a single **fdsinfo**, or **fdsinfo64** structure.

### *IndexPointer*

Specifies the address of a process identifier which indicates the required process table entry. A process identifier of zero selects the first entry in the table. The process identifier is updated to indicate the next entry to be retrieved.

**Note:** The *IndexPointer* does not have to correspond to an existing process, and may in fact correspond to a different process than the one you expect. There is no guarantee that the process slot pointed to by *IndexPointer* will contain the same process between successive calls to **getprocs()** or **getprocs64()**.

*Count* Specifies the number of process table entries requested.

## Return Values

If successful, the **getprocs** subroutine returns the number of process table entries retrieved; if this is less than the number requested, the end of the process table has been reached. A value of 0 is returned when the end of the process table has been reached. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **getprocs** subroutine does not succeed if the following are true:

**EINVAL** The *ProcessSize* or *FileSize* parameters are invalid, or the *IndexPointer* parameter does not point to a valid process identifier, or the *Count* parameter is not greater than zero.

EFAULT

The copy operation to one of the buffers was not successful.

## Related Information

The **getpid** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getpgrp** (“getpid, getpgrp, or getppid Subroutine” on page 407), or **getppid** (“getpid, getpgrp, or getppid Subroutine” on page 407) subroutines, the **getthrds** (“getthrds Subroutine” on page 443) subroutine

The **ps** command.

---

## getproj Subroutine

### Purpose

Retrieves the project definition from the kernel project registry for the requested project name.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
getproj(struct project *, int flag)
```

### Description

The **getproj** subroutine functions similar to the **getprojs** subroutine with the exception that the **getproj** subroutine retrieves the definition only for the project name or number, which is passed as its argument. The *flag* parameter indicates what is passed. The *flag* parameter has the following values:

- PROJ\_NAME — Indicates that the supplied project definition only has the project name. The **getproj** subroutine queries the kernel to obtain a match for the supplied project name and returns the matching entry.
- PROJ\_NUM — Indicates that the supplied project definition only has the project number. The **getproj** subroutine queries the kernel to obtain a match for the supplied project number and returns the matching entry.

Generally, the projects are loaded from the system project definition file or LDAP, or from both. When more than one of these project repositories are used, project name and project ID collisions are possible. These projects are differentiated by the kernel using an origin flag. This origin flag designates the project repository from where the project definition is obtained. If the caller wants to retrieve the project definition that belongs to a specific project repository, the specific origin value should be passed in the flags field of the project structure. Valid project origins values that can be passed are defined in the **sys/aacct.h** file. If the projects are currently loaded from the project repository represented by the origin value, **getproj** returns the specified project if it exists. If the origin value is not passed, the first project reference found in the kernel registry is returned. Regardless of whether the origin is passed or not, **getproj** always returns the project origin flags in the output project structure.

### Parameters

*project*

Pointer holding the project whose information is required.

*flag*

An integer flag that indicates whether the match needs to be performed on the supplied project name or number.

## Security

There are no restrictions. Any user can call this function.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid argument. The <i>flag</i> parameter is not valid or the passed pointer is NULL.
<b>ENOENT</b>	Project not found.

## Related Information

The “addproj Subroutine” on page 31, “chprojattr Subroutine” on page 158, “getprojdb Subroutine,” “getprojs Subroutine” on page 420, rmproj Subroutine.

---

## getprojdb Subroutine

### Purpose

Retrieves the specified project record from the project database.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
getprojdb(void *handle, struct project *project, int flag)
```

### Description

The **getprojdb** subroutine searches the project database associated with the *handle* parameter for the specified project. The project database must be initialized before calling this subroutine. The routines **projdballoc** and **projdbfinit** are provided for this purpose. The *flag* parameter indicates the type of search. The following flags are defined:

- **PROJ\_NAME** — Search by product name. The **getprojdb** subroutine scans the file to obtain a match for the supplied project name and returns the matching entry.
- **PROJ\_NUM** — Search by product number. The **getprojdb** subroutine scans the file to obtain a match for the supplied project number and returns the matching entry.

The entire database is searched. If the specified record is found, the **getprojdb** subroutine stores the relevant project information into the *struct project* buffer, which is passed as an argument to this subroutine. The specified project is then made the current project in the database. If the specified project is not found, the database is reset so that the first project in the database is the current project.

### Parameters

<i>handle</i>	Pointer to the handle allocated for the project database.
<i>project</i>	Pointer holding the project name whose information is required.

*flag* Integer flag indicating what type of information is sent for matching; that is, whether the match needs to be performed by project name or number.

## Security

No restrictions. Any user can call this function.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>ENOENT</b>	Project definition not found.
<b>EINVAL</b>	Invalid arguments if <i>flag</i> is not valid or passed pointer is NULL.

## Related Information

The “addprojdb Subroutine” on page 32, “chprojattrdb Subroutine” on page 159, “getfirstprojdb Subroutine” on page 368, “getnextprojdb Subroutine” on page 396, “getproj Subroutine” on page 418, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, “projdbfree Subroutine” on page 1181, rmprojdb Subroutine.

---

## getprojs Subroutine

### Purpose

Retrieves the project details from the kernel project registry.

### Library

The **libaacct.a** library.

### Syntax

```
<sys/aacct.h>
```

```
getprojs(struct project *, int *)
```

### Description

The **getprojs** subroutine retrieves the specified number of project definitions from the kernel project registry. The number of definitions to be retrieved is passed as an argument to this subroutine, and it is also passed with a buffer of type **struct *project***, where the retrieved project definitions are stored.

When the **getprojs** subroutine is called with a NULL value passed instead of a pointer to a **struct *project***, the **getprojs** subroutine returns the total number of defined projects in the kernel project registry. This number can be used by any subsequent calls to retrieve the project details.

If the integer value passed is smaller than the number of project definitions available, then the project buffer will be filled with as many entries as requested. If the value is greater than the number of available definitions, then the available records are filled in the structure and the integer value is updated with the number of records actually retrieved.

Generally, the projects are loaded from the system project definition file or LDAP, or from both. When more than one of these project repositories are used, project name and project ID collisions are possible. These projects are differentiated by the kernel using an origin flag. This origin flag designates the project repository from where the project definition is obtained. Valid project origins values that can be passed are defined in the `sys/aacct.h` file. The `getproj` subroutine also returns this origin information in the `flags` field of the output project structures.

## Parameters

*pointer*           Points to a project structure where the retrieved data is stored.  
*int*                An integer that indicates the number of elements to be retrieved.

## Security

There are no restrictions. Any user can call this function.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Invalid arguments if passed <i>int</i> pointer is NULL
<b>ENOENT</b>	No projects available.

## Related Information

The “addproj Subroutine” on page 31, “chprojattr Subroutine” on page 158, “getproj Subroutine” on page 418, rmproj Subroutine.

---

## getpw Subroutine

### Purpose

Retrieves a user’s `/etc/passwd` file entry.

### Library

Standard C Library (**libc.a**)

### Syntax

**int** `getpw` (*UserID*, *Buffer*)

**uid\_t** *UserID*  
**char** \**Buffer*

### Description

The `getpw` subroutine opens the `/etc/passwd` file and returns, in the *Buffer* parameter, the `/etc/passwd` file entry of the user specified by the *UserID* parameter.

### Parameters

*Buffer*           Specifies a character buffer large enough to hold any `/etc/passwd` entry.

*UserID* Specifies the ID of the user for which the entry is desired.

## Return Values

The **getpw** subroutine returns:

0 Successful completion  
-1 Not successful.

---

## getpwent, getpwuid, getpwnam, putpwent, setpwent, or endpwent Subroutine

### Purpose

Accesses the basic user information in the user database.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <pwd.h>
struct passwd *getpwent ( )

struct passwd *getpwuid ( UserID)
uid_t UserID;

struct passwd *getpwnam ( Name)
char *Name;

int putpwent ( Password, File)
struct passwd *Password;
FILE *File;

void setpwent ( )
void endpwent ( )
```

### Description

**Attention:** All information generated by the **getpwent**, **getpwnam**, and **getpwuid** subroutines is stored in a static area. Subsequent calls to these subroutines overwrite this static area. To save the information in the static area, applications should copy it.

These subroutines access the basic user attributes.

The **setpwent** subroutine opens the user database if it is not already open. Then, this subroutine sets the cursor to point to the first user entry in the database. The **endpwent** subroutine closes the user database.

The **getpwent**, **getpwnam**, and **getpwuid** subroutines return information about a user. These subroutines do the following:

<b>getpwent</b>	Returns the next user entry in the sequential search.
<b>getpwnam</b>	Returns the first user entry in the database whose name matches the <i>Name</i> parameter.
<b>getpwuid</b>	Returns the first user entry in the database whose ID matches the <i>UserID</i> parameter.

The **putpwent** subroutine writes a password entry into a file in the colon-separated format of the **/etc/passwd** file.

## The passwd Structure

The **getpwent**, **getpwnam**, and **getpwuid** subroutines return a **passwd** structure. The **passwd** structure is defined in the **pwd.h** file and has the following fields:

<code>pw_name</code>	Contains the name of the user name.
<code>pw_passwd</code>	Contains the user's encrypted password. <b>Note:</b> If the password is not stored in the <b>/etc/passwd</b> file and the invoker does not have access to the shadow file that contains passwords, this field contains an undecryptable string, usually an * (asterisk).
<code>pw_uid</code>	Contains the user's ID.
<code>pw_gid</code>	Identifies the user's principal group ID.
<code>pw_gecos</code>	Contains general user information.
<code>pw_dir</code>	Identifies the user's home directory.
<code>pw_shell</code>	Identifies the user's login shell.

**Note:** If Network Information Services (NIS) is enabled on the system, these subroutines attempt to retrieve the information from the NIS authentication server before attempting to retrieve the information locally.

## Parameters

<i>File</i>	Points to an open file whose format is similar to the <b>/etc/passwd</b> file format.
<i>Name</i>	Specifies the user name.
<i>Password</i>	Points to a password structure. This structure contains user attributes.
<i>UserID</i>	Specifies the user ID.

## Security

Files Accessed:

Mode	File
rw	<b>/etc/passwd</b> (write access for the <b>putpwent</b> subroutine only)
r	<b>/etc/security/passwd</b> (if the password is desired)

## Return Values

The **getpwent**, **getpwnam**, and **getpwuid** subroutines return a pointer to a valid password structure if successful. Otherwise, a null pointer is returned.

The **getpwent** subroutine will return a null pointer and an **errno** value of **ENOATTR** when it detects a corrupt entry. To get subsequent entries following the corrupt entry, call the **getpwent** subroutine again.

## Files

<b>/etc/passwd</b>	Contains user IDs and their passwords
--------------------	---------------------------------------

## Related Information

The **getgrent** (“getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371) subroutine, **getgroupattr** (“getgroupattr, IDtogroup, nextgroup, or putgroupattr Subroutine” on page 375) subroutine,

**getuserattr** (“getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454) subroutine, **getuserpw**, **putuserpw**, or **putuserpwhist** (“getuserpw, putuserpw, or putuserpwhist Subroutine” on page 468) subroutine, **setuserdb** subroutine.

List of Security and Auditing Subroutines, Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getrlimit, getrlimit64, setrlimit, setrlimit64, or vlimit Subroutine

### Purpose

Controls maximum system resource consumption.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/time.h>
#include <sys/resource.h>
```

```
int setrlimit( Resource1, RLP)
int Resource1;
struct rlimit *RLP;
```

```
int setrlimit64 ( Resource1, RLP)
int Resource1;
struct rlimit64 *RLP;
```

```
int getrlimit ( Resource1, RLP)
int Resource1;
struct rlimit *RLP;
```

```
int getrlimit64 ( Resource1, RLP)
int Resource1;
struct rlimit64 *RLP;
```

```
#include <sys/vlimit.h>
```

```
vlimit ( Resource2, Value)
int Resource2, Value;
```

### Description

The **getrlimit** subroutine returns the values of limits on system resources used by the current process and its children processes. The **setrlimit** subroutine sets these limits. The **vlimit** subroutine is also supported, but the **getrlimit** subroutine replaces it.

A resource limit is specified as either a soft (current) or hard limit. A calling process can raise or lower its own soft limits, but it cannot raise its soft limits above its hard limits. A calling process must have root user authority to raise a hard limit.

**Note:** The initial values returned by the **getrlimit** subroutine are the ulimit values in effect when the process was started. For maxdata programs the initial soft limit for data is set to the lower of data ulimit value or a value corresponding to the number of data segments reserved for data segments. When a program is executing using the large address-space model, the operating system attempts



to modify the soft limit on data size to match the *maxdata* value. If the *maxdata* value is larger than the current hard limit on data size, either the program will not execute if the **XPG\_SUS\_ENV** environment variable has the value set to ON, or the soft limit will be set to the current hard limit. If the *maxdata* value is smaller than the size of the program's static data, the program will not execute.

The **rlimit** structure specifies the hard and soft limits for a resource, as defined in the **sys/resource.h** file. The **RLIM\_INFINITY** value defines an infinite value for a limit.

When compiled in 32-bit mode, **RLIM\_INFINITY** is a 32-bit value; when compiled in 64-bit mode, it is a 64-bit value. 32-bit routines should use **RLIM64\_INFINITY** when setting 64-bit limits with the **setrlimit64** routine, and recognize this value when returned by **getrlimit64**.

This information is stored as per-process information. This subroutine must be executed directly by the shell if it is to affect all future processes created by the shell.

**Note:** Raising the data limit does not raise the program break value. Use the **brk/sbrk** subroutines to raise the break value. If the proper memory segments are not initialized at program load time, raising your memory limit will not allow access to this memory. Use the **-bmaxdata** flag of the **ld** command to set up these segments at load time.

When compiled in 32-bit mode, the **struct rlimit** values may be returned as **RLIM\_SAVED\_MAX** or **RLIM\_SAVED\_CUR** when the actual resource limit is too large to represent as a 32-bit **rlim\_t**.

These values can be used by library routines which set their own **rlimits** to save off potentially 64-bit **rlimit** values (and prevent them from being truncated by the 32-bit **struct rlimit**). Unless the library routine intends to permanently change the **rlimits**, the **RLIM\_SAVED\_MAX** and **RLIM\_SAVED\_CUR** values can be used to restore the 64-bit **rlimits**.

Application limits may be further constrained by available memory or implementation defined constants such as **OPEN\_MAX** (maximum available open files).

## Parameters

*Resource1*

Can be one of the following values:

### **RLIMIT\_AS**

The maximum size of a process' total available memory, in bytes. This limit is not enforced.

### **RLIMIT\_CORE**

The largest size, in bytes, of a **core** file that can be created. This limit is enforced by the kernel. If the value of the **RLIMIT\_FSIZE** limit is less than the value of the **RLIMIT\_CORE** limit, the system uses the **RLIMIT\_FSIZE** limit value as the soft limit.

### **RLIMIT\_CPU**

The maximum amount of central processing unit (CPU) time, in seconds, to be used by each process. If a process exceeds its soft CPU limit, the kernel will send a **SIGXCPU** signal to the process. After the hard limit is reached, the process will be killed with **SIGXCPU**, even if it handles, blocks, or ignores that signal.

### **RLIMIT\_DATA**

The maximum size, in bytes, of the data region for a process. This limit defines how far a program can extend its break value with the **sbrk** subroutine. This limit is enforced by the kernel. If the **XPG\_SUS\_ENV=ON** environment variable is set in the user's environment before the process is executed and a process attempts to set the limit lower than current usage, the operation fails with **errno** set to **EINVAL**. If the **XPG\_SUS\_ENV** environment variable is not set, the operation fails with **errno** set to **EFAULT**.

### **RLIMIT\_FSIZE**

The largest size, in bytes, of any single file that can be created. When a process attempts to write, truncate, or clear beyond its soft **RLIMIT\_FSIZE** limit, the operation will fail with **errno** set to **EFBIG**. If the environment variable **XPG\_SUS\_ENV=ON** is set in the user's environment before the process is executed, then the **SIGXFSZ** signal is also generated.

### **RLIMIT\_NOFILE**

This is a number one greater than the maximum value that the system may assign to a newly-created descriptor.

### **RLIMIT\_STACK**

The maximum size, in bytes, of the stack region for a process. This limit defines how far a program stack region can be extended. Stack extension is performed automatically by the system. This limit is enforced by the kernel. When the stack limit is reached, the process receives a **SIGSEGV** signal. If this signal is not caught by a handler using the signal stack, the signal ends the process.

### **RLIMIT\_RSS**

The maximum size, in bytes, to which the resident set size of a process can grow. This limit is not enforced by the kernel. A process may exceed its soft limit size without being ended.

*RLP*

Points to the **rlimit** or **rlimit64** structure, which contains the soft (current) and hard limits. For the **getrlimit** subroutine, the requested limits are returned in this structure. For the **setrlimit** subroutine, the desired new limits are specified here.

*Resource2*

The flags for this parameter are defined in the **sys/vlimit.h**, and are mapped to corresponding flags for the **setrlimit** subroutine.

*Value*

Specifies an integer used as a soft-limit parameter to the **vlimit** subroutine.

## Return Values

On successful completion, a return value of 0 is returned, changing or returning the resource limit. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error. If the current limit specified is beyond the hard limit, the **setrlimit** subroutine sets the limit to to max limit and returns successfully.

## Error Codes

The **getrlimit**, **getrlimit64**, **setrlimit**, **setrlimit64**, or **vlimit** subroutine is unsuccessful if one of the following is true:

<b>EFAULT</b>	The address specified for the <i>RLP</i> parameter is not valid.
<b>EINVAL</b>	The <i>Resource1</i> parameter is not a valid resource, or the limit specified in the <i>RLP</i> parameter is invalid.
<b>EPERM</b>	The limit specified to the <b>setrlimit</b> subroutine would have raised the maximum limit value, and the caller does not have root user authority.

## Related Information

The **sigaction**, **sigvec**, or **signal** subroutines, **sigstack** subroutine, **ulimit** subroutine.

---

## getrpcent, getrpcbyname, getrpcbynumber, setrpcent, or endrpcent Subroutine

### Purpose

Accesses the */etc/rpc* file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>

struct rpcent *getrpcent ()
struct rpcent *getrpcbyname ( Name)
char *Name;
struct rpcent *getrpcbynumber ( Number)
int Number;
void setrpcent (StayOpen)
int StayOpen
void endrpcent
```

### Description

**Attention:** Do not use the **getrpcent**, **getrpcbyname**, **getrpcbynumber**, **setrpcent**, or **endrpcent** subroutine in a multithreaded environment.

**Attention:** The information returned by the **getrpcbyname**, and **getrpcbynumber** subroutines is stored in a static area and is overwritten on subsequent calls. Copy the information to save it.

The **getrpcbyname** and **getrpcbynumber** subroutines each return a pointer to an object with the **rpcent** structure. This structure contains the broken-out fields of a line from the */etc/rpc* file. The **getrpcbyname** and **getrpcbynumber** subroutines searches the **rpc** file sequentially from the beginning of the file until it finds a matching RPC program name or number, or until it reaches the end of the file. The **getrpcent** subroutine reads the next line of the file, opening the file if necessary.

The **setrpcent** subroutine opens and rewinds the */etc/rpc* file. If the *StayOpen* parameter does not equal 0, the **rpc** file is not closed after a call to the **getrpcent** subroutine.

The **setrpcent** subroutine rewinds the **rpc** file. The **endrpcent** subroutine closes it.

The **rpc** file contains information about Remote Procedure Call (RPC) programs. The **rpcent** structure is in the **/usr/include/netdb.h** file and contains the following fields:

<code>r_name</code>	Contains the name of the server for an RPC program
<code>r_aliases</code>	Contains an alternate list of names for RPC programs. This list ends with a 0.
<code>r_number</code>	Contains a number associated with an RPC program.

## Parameters

<i>Name</i>	Specifies the name of a server for <b>rpc</b> program.
<i>Number</i>	Specifies the <b>rpc</b> program number for service.
<i>StayOpen</i>	Contains a value used to indicate whether to close the <b>rpc</b> file.

## Return Values

These subroutines return a null pointer when they encounter the end of a file or an error.

## Files

<code>/etc/rpc</code>	Contains information about Remote Procedure Call (RPC) programs.
-----------------------	--

## Related Information

Remote Procedure Call (RPC) for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

---

## getrusage, getrusage64, times, or vtimes Subroutine

### Purpose

Displays information about resource use.

### Libraries

**getrusage, getrusage64, times:** Standard C Library (**libc.a**)

**vtimes:** Berkeley Compatibility Library (**libbsd.a**)

### Syntax

```
#include <sys/times.h>
#include <sys/resource.h>
```

```
int getrusage ( Who, RUsage)
int Who;
struct rusage *RUsage;
```

```
int getrusage64 ( Who, RUsage)
int Who;
struct rusage64 *RUsage;
#include <sys/types.h>
#include <sys/times.h>
```

```
clock_t times ( Buffer)
struct tms *Buffer;
```

```
#include <sys/times.h>

vtimes ( ParentVm, ChildVm)
struct vtimes *ParentVm, ChildVm;
```

## Description

The **getrusage** subroutine displays information about how resources are used by the current process or all completed child processes.

When compiled in 64-bit mode, **rusage** counters are 64 bits. If **getrusage** is compiled in 32-bit mode, **rusage** counters are 32 bits. If the kernel's value of a **usage** counter has exceeded the capacity of the corresponding 32-bit **rusage** value being returned, the **rusage** value is set to INT\_MAX.

The **getrusage64** subroutine can be called to make 64-bit **rusage** counters explicitly available in a 32-bit environment.

In AIX 5.1 and later, 64-bit quantities are also available to 64-bit applications through the **getrusage()** interface in the ru\_utime and ru\_stime fields of **struct rusage**.

The **times** subroutine fills the structure pointed to by the *Buffer* parameter with time-accounting information. All time values reported by the **times** subroutine are measured in terms of the number of clock ticks used. Applications should use **sysconf** (**\_SC\_CLK\_TCK**) to determine the number of clock ticks per second.

The **tms** structure defined in the **/usr/include/sys/times.h** file contains the following fields:

```
time_t  tms_utime;
time_t  tms_stime;
time_t  tms_cutime;
time_t  tms_cstime;
```

This information is read from the calling process as well as from each completed child process for which the calling process executed a **wait** subroutine.

tms_utime	The CPU time used for executing instructions in the user space of the calling process
tms_stime	The CPU time used by the system on behalf of the calling process.
tms_cutime	The sum of the tms_utime and the tms_cutime values for all the child processes.
tms_cstime	The sum of the tms_stime and the tms_cstime values for all the child processes.

**Note:** The system measures time by counting clock interrupts. The precision of the values reported by the **times** subroutine depends on the rate at which the clock interrupts occur.

The **vtimes** subroutine is supported to provide compatibility with earlier programs.

The **vtimes** subroutine returns accounting information for the current process and for the completed child processes of the current process. Either the *ParentVm* parameter, the *ChildVm* parameter, or both may be 0. In that case, only the information for the nonzero pointers is returned.

After a call to the **vtimes** subroutine, each buffer contains information as defined by the contents of the **/usr/include/sys/vtimes.h** file.

## Parameters

*Who* Specifies a value of **RUSAGE\_THREAD**, **RUSAGE\_SELF**, or **RUSAGE\_CHILDREN**.

## *RUsage*

Points to a buffer described in the `/usr/include/sys/resource.h` file. The fields are interpreted as follows:

- ru\_utime**  
The total amount of time running in user mode.
- ru\_stime**  
The total amount of time spent in the system executing on behalf of the processes.
- ru\_maxrss**  
The maximum size, in kilobytes, of the used resident set size.
- ru\_ixrss**  
An integral value indicating the amount of memory used by the text segment that was also shared among other processes. This value is expressed in units of kilobytes \* seconds-of-execution and is calculated by adding the number of shared memory pages in use each time the internal system clock ticks, and then averaging over one-second intervals.
- ru\_idrss**  
An integral value of the amount of unshared memory in the data segment of a process (expressed in units of kilobytes \* seconds-of-execution).
- ru\_minflt**  
The number of page faults serviced without any I/O activity. In this case, I/O activity is avoided by reclaiming a page frame from the list of pages awaiting reallocation.
- ru\_majflt**  
The number of page faults serviced that required I/O activity.
- ru\_nswap**  
The number of times a process was swapped out of main memory.
- ru\_inblock**  
The number of times the file system performed input.
- ru\_oublock**  
The number of times the file system performed output.  
**Note:** The numbers that the `ru_inblock` and `ru_oublock` fields display account for real I/O only; data supplied by the caching mechanism is charged only to the first process to read or write the data.
- ru\_msgsnd**  
The number of IPC messages sent.
- ru\_msgrcv**  
The number of IPC messages received.
- ru\_nsignals**  
The number of signals delivered.
- ru\_nvcsw**  
The number of times a context switch resulted because a process voluntarily gave up the processor before its time slice was completed. This usually occurs while the process waits for availability of a resource.
- ru\_nivcsw**  
The number of times a context switch resulted because a higher priority process ran or because the current process exceeded its time slice.

*Buffer*

Points to a **tms** structure.

*ParentVm*

Points to a **vtimes** structure that contains the accounting information for the current process.

*ChildVm*

Points to a **vtimes** structure that contains the accounting information for the terminated child processes of the current process.

## Return Values

Upon successful completion, the **getrusage** and **getrusage64** subroutines return a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

Upon successful completion, the **times** subroutine returns the elapsed real time in units of ticks, whether profiling is enabled or disabled. This reference time does not change from one call of the **times** subroutine to another. If the **times** subroutine fails, it returns a value of -1 and sets the **errno** global variable to indicate the error.

## Error Codes

The **getrusage** and **getrusage64** subroutines do not run successfully if either of the following is true:

**EINVAL** The *Who* parameter is not a valid value.  
**EFAULT** The address specified for *RUsage* is not valid.

The **times** subroutine does not run successfully if the following is true:

**EFAULT** The address specified by the *buffer* parameter is not valid.

## Related Information

The **gettimer**, **settimer**, **restimer**, **stime**, or **time** (“*gettimer*, *settimer*, *restimer*, *stime*, or *time* Subroutine” on page 446) subroutine, **wait**, **waitpid**, or **wait3** subroutine.

Performance-Related Subroutines in *Performance management*.

---

## getroleattr, nextrole or putroleattr Subroutine

### Purpose

Accesses the role information in the roles database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int getroleattr(Role, Attribute, Value, Type)
char *Role;
char *Attribute;
void *Value;
int Type;
char *nextrole(void)
int putroleattr(Role, Attribute, Value, Type)
char *Role;
char *Attribute;
void *Value;
int Type;
```

### Description

The **getroleattr** subroutine reads a specified attribute from the role database. If the database is not already open, this subroutine does an implicit open for reading.

Similarly, the **putroleattr** subroutine writes a specified attribute into the role database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by the

**putroleattr** subroutine must be explicitly committed by calling the **putroleattr** subroutine with a Type parameter specifying SEC\_COMMIT. Until all the data is committed, only the **getroleattr** subroutine within the process returns written data.

The **nextrole** subroutine returns the next role in a linear search of the role database. The consistency of consecutive searches depends upon the underlying storage-access mechanism and is not guaranteed by this subroutine.

The **setroledb** and **endroledb** subroutines should be used to open and close the role database.

## Parameters

*Attribute* Specifies which attribute is read. The following possible attributes are defined in the **usersec.h** file:

**S\_ROLELIST**

List of roles included by this role. The attribute type is **SEC\_LIST**.

**S\_AUTHORIZATIONS**

List of authorizations included by this role. The attribute type is **SEC\_LIST**.

**S\_GROUPS**

List of groups required for this role. The attribute type is **SEC\_LIST**.

**S\_SCREENs**

List of SMIT screens required for this role. The attribute type is **SEC\_LIST**.

**S\_VISIBILITY**

Number value stating the visibility of the role. The attribute type is **SEC\_INT**.

**S\_MSGCAT**

Message catalog file name. The attribute type is **SEC\_CHAR**.

**S\_MSGNUMBER**

Message number within the catalog. The attribute type is **SEC\_INT**.



<i>Type</i>	Specifies the type of attribute expected. Valid types are defined in the <b>usersec.h</b> file and include: <b>SEC_INT</b> The format of the attribute is an integer. For the <b>getroleattr</b> subroutine, the user should supply a pointer to a defined integer variable. For the <b>putroleattr</b> subroutine, the user should supply an integer. <b>SEC_CHAR</b> The format of the attribute is a null-terminated character string. For the <b>getroleattr</b> subroutine, the user should supply a pointer to a defined character pointer variable. For the <b>putroleattr</b> subroutine, the user should supply a character pointer. <b>SEC_LIST</b> The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series must be an empty (zero character count) string. For the <b>getroleattr</b> subroutine, the user should supply a pointer to a defined character pointer variable. For the <b>putroleattr</b> subroutine, the user should supply a character pointer. <b>SEC_COMMIT</b> For the <b>putroleattr</b> subroutine, this value specified by itself indicates that changes to the named role are to be committed to permanent storage. The <i>Attribute</i> and <i>Value</i> parameters are ignored. If no role is specified, the changes to all modified roles are committed to permanent storage. <b>SEC_DELETE</b> The corresponding attribute is deleted from the database. <b>SEC_NEW</b> Updates the role database file with the new role name when using the <b>putroleattr</b> subroutine.
<i>Value</i>	Specifies a buffer, a pointer to a buffer, or a pointer to a pointer depending on the <i>Attribute</i> and <i>Type</i> parameters. See the <i>Type</i> parameter for more details.

## Return Values

If successful, the **getroleattr** returns 0. Otherwise, a value of -1 is returned and the **errno** global variables is set to indicate the error.

## Error Codes

Possible return codes are:

<b>EACCES</b>	Access permission is denied for the data request.
<b>ENOENT</b>	The specified <i>Role</i> parameter does not exist.
<b>ENOATTR</b>	The specified role attribute does not exist for this role.
<b>EINVAL</b>	The <i>Attribute</i> parameter does not contain one of the defined attributes or null.
<b>EINVAL</b>	The <i>Value</i> parameter does not point to a valid buffer or to valid data for this type of attribute.
<b>EPERM</b>	Operation is not permitted.

## Related Information

The **getuserattr**, **nextusracl**, or **putusraclattr** (“getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454) subroutine, **setroledb**, or **endacldb** subroutine.

---

## gets or fgets Subroutine

### Purpose

Gets a string from a stream.

### Library

Standard I/O Library (**libc.a**)

### Syntax

```
#include <stdio.h>
char *gets ( String )
char *String;
```

```
char *fgets ( String, Number, Stream )
char *String;
int Number;
FILE *Stream;
```

### Description

The **gets** subroutine reads bytes from the standard input stream, **stdin**, into the array pointed to by the *String* parameter. It reads data until it reaches a new-line character or an end-of-file condition. If a new-line character stops the reading process, the **gets** subroutine discards the new-line character and terminates the string with a null character.

The **fgets** subroutine reads bytes from the data pointed to by the *Stream* parameter into the array pointed to by the *String* parameter. The **fgets** subroutine reads data up to the number of bytes specified by the *Number* parameter minus 1, or until it reads a new-line character and transfers that character to the *String* parameter, or until it encounters an end-of-file condition. The **fgets** subroutine then terminates the data string with a null character.

The first successful run of the **fgetc** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **fgets**, **fgetc** (“getc, fgetc, or getwchar Subroutine” on page 477), **fgetws** (“getws or fgets Subroutine” on page 480), **fread** (“fread or fwrite Subroutine” on page 310), **fscanf**, **getc** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **getchar** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **gets** or **scanf** subroutine using a stream that returns data not supplied by a prior call to the **ungetc** or **ungetc** subroutine marks the *st\_atime* field for update.

### Parameters

<i>String</i>	Points to a string to receive bytes.
<i>Stream</i>	Points to the <b>FILE</b> structure of an open file.
<i>Number</i>	Specifies the upper bound on the number of bytes to read.

### Return Values

If the **gets** or **fgets** subroutine encounters the end of the file without reading any bytes, it transfers no bytes to the *String* parameter and returns a null pointer. If a read error occurs, the **gets** or **fgets** subroutine returns a null pointer and sets the **errno** global variable (errors are the same as for the **fgetc** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine). Otherwise, the **gets** or **fgets** subroutine returns the value of the *String* parameter.

**Note:** Depending upon which library routine the application binds to, this subroutine may return **EINTR**. Refer to the **signal** subroutine regarding the **SA\_RESTART** value.

## Related Information

The **feof**, **error**, **clearerr**, or **fileno** (“feof, error, clearerr, or fileno Macro” on page 270) macro, **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fread** (“fread or fwrite Subroutine” on page 310) subroutine, **getc**, **getchar**, **fgetc**, or **getw** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **getwc**, **fgetwc**, or **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **getws** or **fgetws** (“getws or fgetws Subroutine” on page 480) subroutine, **puts** or **fputs** (“puts or fputs Subroutine” on page 1331) subroutine, **putws** or **fputws** (“putws or fputws Subroutine” on page 1341) subroutine, **scanf**, **fscanf**, or **sscanf** subroutine, **ungetc** or **ungetwc** subroutine.

List of String Manipulation Services, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getfsent\_r, getfsspec\_r, getfsfile\_r, getfstype\_r, setfsent\_r, or endfsent\_r Subroutine

### Purpose

Gets information about a file system.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
#include <fstab.h>
```

```
int getfsent_r (FSSent, FSFile, PassNo)
struct fstab * FSSent;
AFILE_t * FSFile;
int * PassNo;
```

```
int getfsspec_r (Special, FSSent, FSFile, PassNo)
const char * Special;
struct fstab *FSSent;
AFILE_t *FSFile;
int *PassNo;
```

```
int getfsfile_r (File, FSSent, FSFile, PassNo)
const char * File;
struct fstab *FSSent;
AFILE_t *FSFile;
int *PassNo;
```

```
int getfstype_r (Type, FSSent, FSFile, PassNo)
const char * Type;
struct fstab *FSSent;
AFILE_t *FSFile;
int *PassNo;
```

```
int setfsent_r (FSFile, PassNo)
AFILE_t * FSFile;
int *PassNo;
```

```
int endfsent_r (FSFile)
AFILE_t *FSFile;
```

## Description

The **getfsent\_r** subroutine reads the next line of the **/etc/filesystems** file, opening it necessary.

The **setfsent\_r** subroutine opens the **filesystems** file and positions to the first record.

The **endfsent\_r** subroutine closes the **filesystems** file.

The **getfsspec\_r** and **getfsfile\_r** subroutines search sequentially from the beginning of the file until a matching special file name or file-system file name is found, or until the end of the file is encountered. The **getfstype\_r** subroutine behaves similarly, matching on the file-system type field.

Programs using this subroutine must link to the **libpthreads.a** library.

## Parameters

<i>FSSent</i>	Points to a structure containing information about the file system. The <i>FSSent</i> parameter must be allocated by the caller. It cannot be a null value.
<i>FSFile</i>	Points to an attribute structure. The <i>FSFile</i> parameter is used to pass values between subroutines.
<i>PassNo</i>	Points to an integer. The <b>setfsent_r</b> subroutine initializes the <i>PassNo</i> parameter.
<i>Special File</i>	Specifies a special file name to search for in the <b>filesystems</b> file.
<i>File</i>	Specifies a file name to search for in the <b>filesystems</b> file.
<i>Type</i>	Specifies a type to search for in the <b>filesystems</b> file.

## Return Values

0	Indicates that the subroutine was successful.
-1	Indicates that the subroutine was not successful.

## Files

<b>/etc/filesystems</b>	Centralizes file-system characteristics.
-------------------------	--

## Related Information

The **getvfsent**, **getvfsbytype**, **getvfsbyname**, **getvfsbyflag**, **setvfsent**, or **endvfsent** (“getvfsent, getvfsbytype, getvfsbyname, getvfsbyflag, setvfsent, or endvfsent Subroutine” on page 476) subroutine.

The **filesystems** file in *AIX 5L Version 5.3 Files Reference*.

List of Multithread Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getsid Subroutine

### Purpose

Returns the session ID of the calling process.

### Library

(**libc.a**)

## Syntax

```
#include <unistd.h>
```

```
pid_t getsid (pid_t pid)
```

## Description

The **getsid** subroutine returns the process group ID of the process that is the session leader of the process specified by *pid*. If *pid* is equal to **pid\_t** subroutine, it specifies the calling process.

## Parameters

*pid*                      A process ID of the process being queried.

## Return Values

Upon successful completion, **getsid** subroutine returns the process group ID of the session leader of the specified process. Otherwise, it returns (**pid\_t**)-1 and set **errno** to indicate the error.

*id*                      The session ID of the requested process.  
-1                      Not successful and the **errno** global variable is set to one of the following error codes.

## Error Codes

**ESRCH**                      There is no process with a process ID equal to *pid*.

**EPERM**                      The process specified by *pid* is not in the same session as the calling process.  
**ESRCH**                      There is no process with a process ID equal to *pid*.

## Related Information

The **exec** (“**exec**: **execl**, **execle**, **execlp**, **execvp**, **execve**, **execvp**, or **exec** Subroutine” on page 236) subroutines, **fork** (“**fork**, **f\_fork**, or **vfork** Subroutine” on page 291) subroutines, **getpid** (“**getpid**, **getpgrp**, or **getppid** Subroutine” on page 407) subroutines, **setpgid** subroutines.

---

## getssys Subroutine

### Purpose

Reads a subsystem record.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <sys/srcobj.h>  
#include <spc.h>
```

```
int getssys( SubsystemName, SRCSubsystem)  
char * SubsystemName;  
struct SRCsubsys * SRCSubsystem;
```

## Description

The **getssys** subroutine reads a subsystem record associated with the specified subsystem and returns the ODM record in the **SRCsubsys** structure.

The **SRCsubsys** structure is defined in the **sys/srcobj.h** file.

## Parameters

<i>SRCSubsystem</i>	Points to the <b>SRCsubsys</b> structure.
<i>SubsystemName</i>	Specifies the name of the subsystem to be read.

## Return Values

Upon successful completion, the **getssys** subroutine returns a value of 0. Otherwise, it returns a value of -1 and the **odmerrno** variable is set to indicate the error, or an SRC error code is returned.

## Error Codes

If the **getssys** subroutine fails, the following is returned:

<b>SRC_NOREC</b>	Subsystem name does not exist.
------------------	--------------------------------

## Files

<b>/etc/objrepos/SRCsubsys</b>	SRC Subsystem Configuration object class.
--------------------------------	---

## Related Information

The **addssys** (“addsys Subroutine” on page 33) subroutine, **delssys** (“delssys Subroutine” on page 212) subroutine, **getsubsvr** (“getsubsvr Subroutine” on page 439) subroutine.

Defining Your Subsystem to the SRC, List of SRC Subroutines, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getsubopt Subroutine

### Purpose

Parse suboptions from a string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
int getsubopt (char **optionp,  
char * const * tokens,  
char ** valuep)
```

### Description

The **getsubopt** subroutine parses suboptions in a flag parameter that were initially parsed by the **getopt** subroutine. These suboptions are separated by commas and may consist of either a single token, or a

token-value pair separated by an equal sign. Because commas delimit suboptions in the option string, they are not allowed to be part of the suboption or the value of a suboption. Similarly, because the equal sign separates a token from its value, a token must not contain an equal sign.

The **getsubopt** subroutine takes the address of a pointer to the option string, a vector of possible tokens, and the address of a value string pointer. It returns the index of the token that matched the suboption in the input string or -1 if there was no match. If the option string at *\*optionp* contains only one suboption, the **getsubopt** subroutine updates *\*optionp* to point to the start of the next suboption. If the suboption has an associated value, the **getsubopt** subroutine updates *\*valuep* to point to the value's first character. Otherwise it sets *\*valuep* to a NULL pointer.

The token vector is organized as a series of pointers to strings. The end of the token vector is identified by a NULL pointer.

When the **getsubopt** subroutine returns, if *\*valuep* is not a NULL pointer then the suboption processed included a value. The calling program may use this information to determine if the presence or lack of a value for this suboption is an error.

Additionally, when the **getsubopt** subroutine fails to match the suboption with the tokens in the *tokens* array, the calling program should decide if this is an error, or if the unrecognized option should be passed on to another program.

## Return Values

The **getsubopt** subroutine returns the index of the matched token string, or -1 if no token strings were matched.

## Related Information

The **getopt** ("getopt Subroutine" on page 397) subroutine.

---

## getsubsvr Subroutine

### Purpose

Reads a subsystem record.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <sys/srcobj.h>
#include <spc.h>
```

```
int getsubsvr( SubserverName, SRCSubserver)
char *SubserverName;
struct SRCSubsvr *SRCSubserver;
```

### Description

The **getsubsvr** subroutine reads a subsystem record associated with the specified subserver and returns the ODM record in the **SRCsubsvr** structure.

The **SRCsubsvr** structure is defined in the **sys/srcobj.h** file and includes the following fields:

```
char          sub_type[30];
```

```
char      subsysname[30];
short     sub_code;
```

## Parameters

*SRCSubserver* Points to the **SRCsubsvr** structure.  
*SubserverName* Specifies the subserver to be read.

## Return Values

Upon successful completion, the **getsubsvr** subroutine returns a value of 0. Otherwise, it returns a value of -1 and the **odmerrno** variable is set to indicate the error, or an SRC error code is returned.

## Error Codes

If the **getsubsvr** subroutine fails, the following is returned:

**SRC\_NOREC** The specified **SRCsubsvr** record does not exist.

## Files

*/etc/objrepos/SRCsubsvr* SRC Subserver Configuration object class.

## Related Information

The **getssys** (“getssys Subroutine” on page 437) subroutine.

Defining Your Subsystem to the SRC, List of SRC Subroutines, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## gettcbatr or puttcbatr Subroutine

### Purpose

Accesses the TCB information in the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int gettcbatr (Entry, Attribute, Value, Type)
```

```
char * Entry;
```

```
char * Attribute;
```

```
void * Value;
```

```
int Type;
```

```
int puttcbatr (Entry, Attribute, Value, Type)
```

```
char *Entry;
```

```
char *Attribute;
```

```
void *Value;
```

```
int Type;
```



## Description

These subroutines access Trusted Computing Base (TCB) information.

The **gettcbatr** subroutine reads a specified attribute from the `tcbck` database. If the database is not already open, the subroutine will do an implicit open for reading.

Similarly, the **puttcbatr** subroutine writes a specified attribute into the `tcbck` database. If the database is not already open, the subroutine does an implicit open for reading and writing. Data changed by **puttcbatr** must be explicitly committed by calling the **puttcbatr** subroutine with a *Type* parameter specifying the **SEC\_COMMIT** value. Until the data is committed, only **get** subroutine calls within the process will return the written data.

New entries in the `tcbck` databases must first be created by invoking **puttcbatr** with the **SEC\_NEW** type.

The `tcbck` database usually defines all the files and programs that are part of the TCB, but the root user or a member of the security group can choose to define only those files considered to be security-relevant.

## Parameters

<i>Attribute</i>	Specifies which attribute is read. The following possible values are defined in the <b>sysck.h</b> file: <b>S_ACL</b> The access control list for the file. Type: <b>SEC_CHAR</b> . <b>S_CHECKSUM</b> The checksum of the file. Type: <b>SEC_CHAR</b> . <b>S_CLASS</b> The logical group of the file. The attribute type is <b>SEC_LIST</b> . <b>S_GROUP</b> The file group. The attribute type is <b>SEC_CHAR</b> . <b>S_LINKS</b> The hard links to this file. Type: <b>SEC_LIST</b> . <b>S_MODE</b> The File mode. Type: <b>SEC_CHAR</b> . <b>S_OWNER</b> The file owner. Type: <b>SEC_CHAR</b> . <b>S_PROGRAM</b> The associated checking program for the file. Type: <b>SEC_CHAR</b> . <b>S_SIZE</b> The size of the file in bytes. Type: <b>SEC_LONG</b> . <b>S_SOURCE</b> The source for the file. Type: <b>SEC_CHAR</b> . <b>S_SYMLINKS</b> The symbolic links to the file. Type: <b>SEC_LIST</b> . <b>S_TARGET</b> The target file (if file is a symbolic link). Type: <b>SEC_CHAR</b> . <b>S_TCB</b> The Trusted Computer Base. The attribute type is <b>SEC_BOOL</b> . <b>S_TYPE</b> The type of file. The attribute type is <b>SEC_CHAR</b> .  Additional user-defined attributes may be used and will be stored in the format specified by the <i>Type</i> parameter.
<i>Entry</i>	Specifies the name of the file for which an attribute is to be read or written.

<i>Type</i>	<p>Specifies the type of attribute expected. Valid values are defined in the <b>usersec.h</b> file and include:</p> <p><b>SEC_BOOL</b> A pointer to an integer (<b>int *</b>) that has been cast to a null pointer.</p> <p><b>SEC_CHAR</b> The format of the attribute is a null-terminated character string.</p> <p><b>SEC_LIST</b> The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series is terminated by two successive null characters.</p> <p><b>SEC_LONG</b> The format of the attribute is a 32-bit integer.</p>
<i>Value</i>	<p>Specifies the address of a pointer for the <b>gettcattr</b> subroutine. The <b>gettcattr</b> subroutine will return the address of a buffer in the pointer. For the <b>puttcattr</b> subroutine, the <i>Value</i> parameter specifies the address of a buffer in which the attribute is stored. See the <i>Type</i> parameter for more details.</p>

## Security

Files Accessed:

Mode	File
rw	<b>/etc/security/sysck.cfg</b> (write access for <b>puttcattr</b> )

## Return Values

The **gettcattr** and **puttcattr** subroutines, when successfully completed, return a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

**Note:** These subroutines return errors from other subroutines.

These subroutines fail if the following is true:

<b>EACCES</b>	Access permission is denied for the data request.
---------------	---

The **gettcattr** and **puttcattr** subroutines fail if one or more of the following are true:

<b>EINVAL</b>	The <i>Value</i> parameter does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors may not be detected.
<b>EINVAL</b>	The <i>Entry</i> parameter is null or contains a pointer to a null string.
<b>EINVAL</b>	The <i>Type</i> parameter contains more than one of the <b>SEC_BOOL</b> , <b>SEC_CHAR</b> , <b>SEC_LIST</b> , or <b>SEC_LONG</b> attributes.
<b>EINVAL</b>	The <i>Type</i> parameter specifies that an individual attribute is to be committed, and the <i>Entry</i> parameter is null.
<b>ENOENT</b>	The specified <i>Entry</i> parameter does not exist or the attribute is not defined for this entry.
<b>EPERM</b>	Operation is not permitted.

## Related Information

The **getuserattr** (“getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454) subroutine, **getuserpw** (“getuserpw, putuserpw, or putuserpw hist Subroutine” on page 468) subroutine, **setpwdb** subroutine, **setuserdb** subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getthrds Subroutine

### Purpose

Gets kernel thread table entries.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <procinfo.h>
#include <sys/types.h>
```

```
int
getthrds ( ProcessIdentifier, ThreadBuffer, ThreadSize, IndexPointer, Count)
pid_t ProcessIdentifier;
struct thrdsinfo *ThreadBuffer;
or struct thrdsinfo64 *ThreadBuffer;
int ThreadSize;
tid_t *IndexPointer;
int Count;
```

```
int
getthrds64 ( ProcessIdentifier, ThreadBuffer, ThreadSize, IndexPointer, Count)
pid_t ProcessIdentifier;
struct thrdentry64 *ThreadBuffer;
int ThreadSize;
tid64_t *IndexPointer;
int Count;
```

### Description

The **getthrds** subroutine returns information about kernel threads, including kernel thread table information defined by the **thrdsinfo** or **thrdsinfo64** structure.

The **getthrds** subroutine retrieves up to *Count* kernel thread table entries, starting with the entry corresponding to the thread identifier indicated by *IndexPointer*, and places them in the array of **thrdsinfo** or **thrdsinfo64**, or **thrdentry64** structures indicated by the *ThreadBuffer* parameter.

On return, the kernel thread identifier referenced by *IndexPointer* is updated to indicate the next kernel thread table entry to be retrieved. The **getthrds** subroutine returns the number of kernel thread table entries retrieved.

If the *ProcessIdentifier* parameter indicates a process identifier, only kernel threads belonging to that process are considered. If this parameter is set to -1, all kernel threads are considered.

The **getthrds** subroutine is normally called repeatedly in a loop, starting with a kernel thread identifier of zero, and looping until the return value is less than *Count*, indicating that there are no more entries to retrieve.

1. Do not use information from the **procsinfo** structure (see the **getprocs** (“getprocs Subroutine” on page 415) subroutine) to determine the value of the *Count* parameter; a process may create or destroy kernel threads in the interval between a call to **getprocs** and a subsequent call to **getthrds**.
2. The kernel thread table may change while the **getthrds** subroutine is accessing it. Returned entries will always be consistent, but since kernel threads can be created or destroyed while the **getthrds** subroutine is running, there is no guarantee that retrieved entries will still exist, or that all existing kernel threads have been retrieved.

When used in 32-bit mode, limits larger than can be represented in 32 bits are truncated to RLIM\_INFINITY. Large values are truncated to INT\_MAX. 64-bit applications are required to use **getthrds64()** and **struct thrdentry64**. Note that **struct thrdentry64** contains the same information as **struct thrdsinfo64** with the only difference being support for the 64-bit tid\_t and the 256-bit sigset\_t. Application developers are also encouraged to use **getthrds64()** in 32-bit applications to obtain 64-bit thread information as this interface provides the new, larger types. The **getthrds()** interface will still be supported for 32-bit applications using **struct thrdsinfo** or **struct thrdsinfo64**, but will not be available to 64-bit applications.

## Parameters

### *ProcessIdentifier*

Specifies the process identifier of the process whose kernel threads are to be retrieved. If this parameter is set to -1, all kernel threads in the kernel thread table are retrieved.

### *ThreadBuffer*

Specifies the starting address of an array of **thrdsinfo** or **thrdsinfo64**, or **thrdentry64** structures which will be filled in with kernel thread table entries. If a value of **NULL** is passed for this parameter, the **getthrds** subroutine scans the kernel thread table and sets return values as normal, but no kernel thread table entries are retrieved.

### *ThreadSize*

Specifies the size of a single **thrdsinfo**, **thrdsinfo64**, or **thrdentry64** structure.

### *IndexPointer*

Specifies the address of a kernel thread identifier which indicates the required kernel thread table entry (this does not have to correspond to an existing kernel thread). A kernel thread identifier of zero selects the first entry in the table. The kernel thread identifier is updated to indicate the next entry to be retrieved.

*Count* Specifies the number of kernel thread table entries requested.

## Return Value

If successful, the **getthrds** subroutine returns the number of kernel thread table entries retrieved; if this is less than the number requested, the end of the kernel thread table has been reached. A value of 0 is returned when the end of the kernel thread table has been reached. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **getthrds** subroutine fails if the following are true:

<b>EINVAL</b>	The <i>ThreadSize</i> is invalid, or the <i>IndexPointer</i> parameter does not point to a valid kernel thread identifier, or the <i>Count</i> parameter is not greater than zero.
<b>ESRCH</b>	The process specified by the <i>ProcessIdentifier</i> parameter does not exist.
<b>EFAULT</b>	The copy operation to one of the buffers failed.

## Related Information

The **getpid** (“getpid, getpgrp, or getppid Subroutine” on page 407), **getpgrp** (“getpid, getpgrp, or getppid Subroutine” on page 407), or **getppid** (“getpid, getpgrp, or getppid Subroutine” on page 407) subroutines, the **getprocs** (“getprocs Subroutine” on page 415) subroutine.

The **ps** command.

---

## gettimeofday, settimeofday, or ftime Subroutine

### Purpose

Displays, gets and sets date and time.

### Libraries

**gettimeofday**, **settimeofday**: Standard C Library (**libc.a**)

**ftime**: Berkeley Compatibility Library (**libbsd.a**)

### Syntax

```
#include <sys/time.h>
int gettimeofday ( Tp, Tzp)
struct timeval *Tp;
void *Tzp;
int settimeofday (Tp, Tzp)
struct timeval *Tp;
struct timezone *Tzp;
#include <sys/types.h>
#include <sys/timeb.h>
int ftime (Tp)
struct timeb *Tp;
```

### Description

Current Greenwich time and the current time zone are displayed with the **gettimeofday** subroutine, and set with the **settimeofday** subroutine. The time is expressed in seconds and microseconds since midnight (0 hour), January 1, 1970. The resolution of the system clock is hardware-dependent, and the time may be updated either continuously or in “ticks.” If the *Tzp* parameter has a value of 0, the time zone information is not returned or set.

If a recent **adjtime** subroutine call is causing the clock to be adjusted backwards, it is possible that two closely spaced **gettimeofday** calls will observe that time has moved backwards. You can set the **GETTOD\_ADJ\_MONOTONIC** environment value to cause the returned value to never decrease. After this environment variable is set, the returned value briefly remains constant as necessary to always report a nondecreasing time of day. This extra processing adds significant pathlength to **gettimeofday**. Although any setting of this environment variable requires this extra processing, setting it to 1 is recommended for future compatibility.

The *Tp* parameter returns a pointer to a **timeval** structure that contains the time since the epoch began in seconds and microseconds.

The **timezone** structure indicates both the local time zone (measured in minutes of time westward from Greenwich) and a flag that, if nonzero, indicates that daylight saving time applies locally during the appropriate part of the year.

In addition to the difference in timer granularity, the **timezone** structure distinguishes these subroutines from the POSIX **gettimer** and **settimer** subroutines, which deal strictly with Greenwich Mean Time.

The **ftime** subroutine fills in a structure pointed to by its argument, as defined by **<sys/timeb.h>**. The structure contains the time in seconds since 00:00:00 UTC (Coordinated Universal Time), January 1, 1970, up to 1000 milliseconds of more-precise interval, the local timezone (measured in minutes of time westward from UTC), and a flag that, if nonzero, indicates that Daylight Saving time is in effect, and the values stored in the **timeb** structure have been adjusted accordingly.

## Parameters

*Tp* Pointer to a **timeval** structure, defined in the **sys/time.h** file.  
*Tzp* Pointer to a **timezone** structure, defined in the **sys/time.h** file.

## Return Values

If the subroutine succeeds, a value of 0 is returned. If an error occurs, a value of -1 is returned and **errno** is set to indicate the error.

## Error Codes

If the **settimeofday** subroutine is unsuccessful, the **errno** value is set to **EPERM** to indicate that the process's effective user ID does not have root user authority.

No errors are defined for the **gettimeofday** or **ftime** subroutine.

---

## gettimer, settimer, restimer, stime, or time Subroutine

### Purpose

Gets or sets the current value for the specified systemwide timer.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/time.h>
#include <sys/types.h>
```

```
int gettimer( TimerType, Value)
timer_t TimerType;
struct timestruc_t * Value;
```

```
#include <sys/timers.h>
#include <sys/types.h>
```

```
int gettimer( TimerType, Value)
timer_t TimerType;
struct itimerspec * Value;
```

```
int settimer(TimerType, TimePointer)
int TimerType;
const struct timestruc_t *TimePointer;
```

```

int restimer(TimerType, Resolution, MaximumValue)
int TimerType;
struct timestruc_t *Resolution, *MaximumValue;

int stime(Tp)
long *Tp;
#include <sys/types.h>
time_t time(Tp)
time_t *Tp;

```

## Description

The **settimer** subroutine is used to set the current value of the *TimePointer* parameter for the systemwide timer, specified by the *TimerType* parameter.

When the **gettimer** subroutine is used with the function prototype in **sys/timers.h**, then except for the parameters, the **gettimer** subroutine is identical to the **getinterval** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) subroutine. Use of the **getinterval** subroutine is recommended, unless the **gettimer** subroutine is required for a standards-conformant application. The description and semantics of the **gettimer** subroutine are subject to change between releases, pending changes in the draft standard upon which the current **gettimer** subroutine description is based.

When the **gettimer** subroutine is used with the function prototype in **/sys/timers.h**, the **gettimer** subroutine returns an **itimerspec** structure to the pointer specified by the *Value* parameter. The **it\_value** member of the **itimerspec** structure represents the amount of time in the current interval before the timer (specified by the *TimerType* parameter) expires, or a zero interval if the timer is disabled. The members of the pointer specified by the *Value* parameter are subject to the resolution of the timer.

When the **gettimer** subroutine is used with the function prototype in **sys/time.h**, the **gettimer** subroutine returns a **timestruc** structure to the pointer specified by the *Value* parameter. This structure holds the current value of the system wide timer specified by the *Value* parameter.

The resolution of any timer can be obtained by the **restimer** subroutine. The *Resolution* parameter represents the resolution of the specified timer. The *MaximumValue* parameter represents the maximum possible timer value. The value of these parameters are the resolution accepted by the **settimer** subroutine.

**Note:** If a nonprivileged user attempts to submit a fine granularity timer (that is, a timer request of less than 10 milliseconds), the timer request is raised to 10 milliseconds.

The **time** subroutine returns the time in seconds since the Epoch (that is, 00:00:00 GMT, January 1, 1970). The *Tp* parameter points to an area where the return value is also stored. If the *Tp* parameter is a null pointer, no value is stored.

The **stime** subroutine is implemented to provide compatibility with older AIX, AT&T System V, and BSD systems. It calls the **settimer** subroutine using the **TIMEOFDAY** timer.

## Parameters

*Value* Points to a structure of type **itimerspec**.

<i>TimerType</i>	Specifies the systemwide timer: <b>TIMEOFDAY</b> (POSIX system clock timer) This timer represents the time-of-day clock for the system. For this timer, the values returned by the <b>gettimer</b> subroutine and specified by the <b>settimer</b> subroutine represent the amount of time since 00:00:00 GMT, January 1, 1970.
<i>TimePointer</i>	Points to a structure of type <b>struct timestruc_t</b> .
<i>Resolution</i>	The resolution of a specified timer.
<i>MaximumValue</i>	The maximum possible timer value.
<i>Tp</i>	Points to a structure containing the time in seconds.

## Return Values

The **gettimer**, **settimer**, **restimer**, and **stime** subroutines return a value of 0 (zero) if the call is successful. A return value of -1 indicates an error occurred, and **errno** is set.

The **time** subroutine returns the value of time in seconds since Epoch. Otherwise, a value of  $((\mathbf{time\_t}) - 1)$  is returned and the **errno** global variable is set to indicate the error.

## Error Codes

If an error occurs in the **gettimer**, **settimer**, **restimer**, or **stime** subroutine, a return value of - 1 is received and the **errno** global variable is set to one of the following error codes:

<b>EINVAL</b>	The <i>TimerType</i> parameter does not specify a known systemwide timer, or the <i>TimePointer</i> parameter of the <b>settimer</b> subroutine is outside the range for the specified systemwide timer.
<b>EFAULT</b>	A parameter address referenced memory that was not valid.
<b>EIO</b>	An error occurred while accessing the timer device.
<b>EPERM</b>	The requesting process does not have the appropriate privilege to set the specified timer.

If the **time** subroutine is unsuccessful, a return value of -1 is received and the **errno** global variable is set to the following:

<b>EFAULT</b>	A parameter address referenced memory that was not valid.
---------------	---

## Related Information

The **asctime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine, **clock** (“clock Subroutine” on page 169) subroutine, **ctime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine, **difftime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine, **getinterval** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) subroutine, **gmtime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine, **localtime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine, **mktime** (“ctime, localtime, gmtime, mktime, difftime, asctime, or tzset Subroutine” on page 200) subroutine, **strptime** subroutine, **strptime** subroutine, **utime** subroutine.

Time data manipulation services in *Operating system and device management*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.



---

## gettimerid Subroutine

### Purpose

Allocates a per-process interval timer.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/time.h>
#include <sys/events.h>
```

```
timer_t gettimerid( TimerType, NotifyType)
int TimerType;
int NotifyType;
```

### Description

The **gettimerid** subroutine is used to allocate a per-process interval timer based on the timer with the given timer type. The unique ID is used to identify the interval timer in interval timer requests. (See **getinterval** subroutine). The particular timer type, the *TimerType* parameter, is defined in the **sys/time.h** file and can identify either a systemwide timer or a per-process timer. The mechanism by which the process is to be notified of the expiration of the timer event is the *NotifyType* parameter, which is defined in the **sys/events.h** file.

The *TimerType* parameter represents one of the following timer types:

<b>TIMEOFDAY</b>	(POSIX system clock timer) This timer represents the time-of-day clock for the system. For this timer, the values returned by the <b>gettimer</b> subroutine and specified by the <b>settimer</b> subroutine represent the amount of time since 00:00:00 GMT, January 1, 1970, in nanoseconds.
<b>TIMERID_ALARM</b>	(Alarm timer) This timer schedules the delivery of a <b>SIGALRM</b> signal at a timer specified in the call to the <b>settimer</b> subroutine.
<b>TIMERID_REAL</b>	(Real-time timer) The real-time timer decrements in real time. A <b>SIGALRM</b> signal is delivered when this timer expires.
<b>TIMERID_VIRTUAL</b>	(Virtual timer) The virtual timer decrements in process virtual time. It runs only when the process is executing in user mode. A <b>SIGVTALRM</b> signal is delivered when it expires.
<b>TIMERID_PROF</b>	(Profiling timer) The profiling timer decrements both when running in user mode and when the system is running for the process. It is designed to be used by processes to profile their execution statistically. A <b>SIGPROF</b> signal is delivered when the profiling timer expires.

Interval timers with a notification value of **DELIVERY\_SIGNAL** are inherited across an **exec** subroutine.

### Parameters

<i>NotifyType</i>	Notifies the process of the expiration of the timer event.
<i>TimerType</i>	Identifies either a systemwide timer or a per-process timer.

## Return Values

If the **gettimerid** subroutine succeeds, it returns a **timer\_t** structure that can be passed to the per-process interval timer subroutines, such as the **getinterval** subroutine. If an error occurs, the value -1 is returned and **errno** is set.

## Error Codes

If the **gettimerid** subroutine fails, the value -1 is returned and **errno** is set to one of the following error codes:

<b>EAGAIN</b>	The calling process has already allocated all of the interval timers associated with the specified timer type for this implementation.
<b>EINVAL</b>	The specified timer type is not defined.

## Related Information

The **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exec Subroutine” on page 236) subroutine, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **getinterval**, **incinterval**, **absinterval**, **resinc**, or **resabs** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) subroutine, **gettimer**, **settimer**, or **restimer** (“gettimer, settimer, restimer, stime, or time Subroutine” on page 446) subroutine, **reltimerid** subroutine.

Time data manipulation services in *Operating system and device management*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getttyent, getttynam, setttyent, or endtttyent Subroutine

### Purpose

Gets a tty description file entry.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <tttyent.h>
```

```
struct tttyent *getttyent()  
struct tttyent *getttynam( Name )  
char *Name;  
void setttyent()  
void endtttyent()
```

### Description

**Attention:** Do not use the **getttyent**, **getttynam**, **setttyent**, or **endtttyent** subroutine in a multithreaded environment.

The **getttyent** and **getttynam** subroutines each return a pointer to an object with the **tttyent** structure. This structure contains the broken-out fields of a line from the tty description file. The **tttyent** structure is in the **/usr/include/sys/tttyent.h** file and contains the following fields:

tty_name	The name of the character special file in the <b>/dev</b> directory. The character special file must reside in the <b>/dev</b> directory.
ty_getty	The command that is called by the <b>init</b> process to initialize tty line characteristics. This is usually the <b>getty</b> command, but any arbitrary command can be used. A typical use is to initiate a terminal emulator in a window system.
ty_type	The name of the default terminal type connected to this tty line. This is typically a name from the <b>termcap</b> database. The <b>TERM</b> environment variable is initialized with this name by the <b>getty</b> or <b>login</b> command.
ty_status	A mask of bit fields that indicate various actions to be allowed on this tty line. The following is a description of each flag:  <b>TTY_ON</b> Enables logins (that is, the <b>init</b> process starts the specified <b>getty</b> command on this entry).  <b>TTY_SECURE</b> Allows a user with root user authority to log in to this terminal. The <b>TTY_ON</b> flag must be included.
ty_window	The command to execute for a window system associated with the line. The window system is started before the command specified in the ty_getty field is executed. If none is specified, this is null.
ty_comment	The trailing comment field. A leading delimiter and white space is removed.

The **gettyent** subroutine reads the next line from the tty file, opening the file if necessary. The **settyent** subroutine rewinds the file. The **endttyent** subroutine closes it.

The **gettynam** subroutine searches from the beginning of the file until a matching name (specified by the *Name* parameter) is found (or until the EOF is encountered).

## Parameters

*Name* Specifies the name of a tty description file.

## Return Values

These subroutines return a null pointer when they encounter an EOF (end-of-file) character or an error.

## Files

<b>/usr/lib/libodm.a</b>	Specifies the ODM (Object Data Manager) library.
<b>/usr/lib/libcfg.a</b>	Archives device configuration subroutines.
<b>/etc/termcap</b>	Defines terminal capabilities.

## Related Information

The **ttyslot** subroutine.

The **getty** command, **init** command, **login** command.

List of Files and Directories Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getuid, geteuid, or getuidx Subroutine

### Purpose

Gets the real or effective user ID of the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <unistd.h>
uid_t getuid(void)
uid_t geteuid(void)
#include <id.h>
uid_t getuidx (int type);
```

### Description

The **getuid** subroutine returns the real user ID of the current process. The **geteuid** subroutine returns the effective user ID of the current process.

The **getuidx** subroutine returns the user ID indicated by the *type* parameter of the calling process.

These subroutines are part of Base Operating System (BOS) Runtime.

### Return Values

The **getuid**, **geteuid** and **getuidx** subroutines return the corresponding user ID. The **getuid** and **geteuid** subroutines always succeed.

The **getuidx** subroutine will return -1 and set the global **errno** variable to **EINVAL** if the *type* parameter is not one of **ID\_REAL**, **ID\_EFFECTIVE**, **ID\_SAVED** or **ID\_LOGIN**.

### Parameters

*type* Specifies the user ID to get. Must be one of **ID\_REAL** (real user ID), **ID\_EFFECTIVE** (effective user ID), **ID\_SAVED** (saved set-user ID) or **ID\_LOGIN** (login user ID).

### Error Codes

If the **getuidx** subroutine fails the following is returned:

**EINVAL** Indicates the value of the *type* parameter is invalid.

### Related Information

The **setuid** subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getuserinfo Subroutine

### Purpose

Finds a value associated with a user.

### Library

Standard C Library (**libc.a**)

### Syntax

```
char *getuserinfo ( Name)
char *Name;
```

### Description

The **getuserinfo** subroutine finds a value associated with a user. This subroutine searches a user information buffer for a string of the form *Name=Value* and returns a pointer to the *Value* substring if the *Name* value is found. A null value is returned if the *Name* value is not found.

The **INuibp** global variable points to the user information buffer:

```
extern char *INuibp;
```

This variable is initialized to a null value.

If the **INuibp** global variable is null when the **getuserinfo** subroutine is called, the **usrinfo** subroutine is called to read user information from the kernel into a local buffer. The **INUuibp** is set to the address of the local buffer. If the **INuibp** external variable is not set, the **usrinfo** subroutine is automatically called the first time the **getuserinfo** subroutine is called.

### Parameter

*Name*                      Specifies a user name.

### Related Information

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getuserinfo Subroutine

### Purpose

Finds a value associated with a user.

### Library

Standard C Library (**libc.a**)

### Syntax

```
char *getuserinfo ( Name)
char *Name;
```

## Description

The **getuinfox** subroutine finds a value associated with a user. This subroutine searches a privileged kernel buffer for a string of the form *Name=Value* and returns a pointer to the *Value* substring if the *Name* value is found. A Null value is returned if the *Name* value is not found. The caller is responsible for freeing the memory returned by the **getuinfox** subroutine.

## Parameters

*Name*                      Specifies a name.

## Return Values

Upon success, the **getuinfox** subroutine returns a pointer to the *Value* substring.

## Error Codes

A Null value is returned if the *Name* value is not found.

---

## getuserattr, IDtouser, nextuser, or putuserattr Subroutine

### Purpose

Accesses the user information in the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int getuserattr (User, Attribute, Value, Type)
char * User;
char * Attribute;
void * Value;
int Type;
```

```
char *IDtouser( UID)
uid_t UID;
```

```
char *nextuser ( Mode, Argument)
int Mode, Argument;
int putuserattr (User, Attribute, Value, Type)
char *User;
char *Attribute;
void *Value;
int Type;
```

### Description

**Attention:** These subroutines and the **setpwent** and **setgrent** subroutines should not be used simultaneously. The results can be unpredictable.

These subroutines access user information. Because of their greater granularity and extensibility, you should use them instead of the **getpwent** routines.

The **getuserattr** subroutine reads a specified attribute from the user database. If the database is not already open, this subroutine does an implicit open for reading. A call to the **getuserattr** subroutine for every new user verifies that the user exists.

Similarly, the **putuserattr** subroutine writes a specified attribute into the user database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by the **putuserattr** subroutine must be explicitly committed by calling the **putuserattr** subroutine with a *Type* parameter specifying **SEC\_COMMIT**. Until all the data is committed, only these subroutines within the process return written data.

New entries in the user and group databases must first be created by invoking **putuserattr** with the **SEC\_NEW** type.

The **IDtouser** subroutine translates a user ID into a user name.

The **nextuser** subroutine returns the next user in a linear search of the user database. The consistency of consecutive searches depends upon the underlying storage-access mechanism and is not guaranteed by this subroutine.

The **setuserdb** and **enduserdb** subroutines should be used to open and close the user database.

The **enduserdb** subroutine frees all memory allocated by the **getuserattr** subroutine.

## Parameters

### *Argument*

Presently unused and must be specified as null.

### *Attribute*

Specifies which attribute is read. The following possible attributes are defined in the **usersec.h** file:

#### **S\_CORECOMP**

Core compression status. The attribute type is **SEC\_CHAR**.

#### **S\_COREPATH**

Core path specification status. The attribute type is **SEC\_CHAR**.

#### **S\_COREPNAME**

Core path specification location. The attribute type is **SEC\_CHAR**.

#### **S\_CORENAMING**

Core naming status. The attribute type is **SEC\_CHAR**.

**S\_ID** User ID. The attribute type is **SEC\_INT**.

#### **S\_PGID**

Principle group ID. The attribute type is **SEC\_INT**.

#### **S\_PGRP**

Principle group name. The attribute type is **SEC\_CHAR**.

#### **S\_GROUPS**

Groups to which the user belongs. The attribute type is **SEC\_LIST**.

#### **S\_ADMGROUPS**

Groups for which the user is an administrator. The attribute type is **SEC\_LIST**.

#### **S\_ADMIN**

Administrative status of a user. The attribute type is **SEC\_BOOL**.

#### **S\_AUDITCLASSES**

Audit classes to which the user belongs. The attribute type is **SEC\_LIST**.

**S\_AUTHSYSTEM**

Defines the user's authentication method. The attribute type is **SEC\_CHAR**.

**S\_HOME**

Home directory. The attribute type is **SEC\_CHAR**.

**S\_SHELL**

Initial program run by a user. The attribute type is **SEC\_CHAR**.

**S\_GECOS**

Personal information for a user. The attribute type is **SEC\_CHAR**.

**S\_USRENV**

User-state environment variables. The attribute type is **SEC\_LIST**.

**S\_SYSENV**

Protected-state environment variables. The attribute type is **SEC\_LIST**.

**S\_LOGINCHK**

Specifies whether the user account can be used for local logins. The attribute type is **SEC\_BOOL**.

**S\_HISTEXPIRE**

Defines the period of time (in weeks) that a user cannot reuse a password. The attribute type is **SEC\_INT**.

**S\_HISTSIZE**

Specifies the number of previous passwords that the user cannot reuse. The attribute type is **SEC\_INT**.

**S\_MAXREPEAT**

Defines the maximum number of times a user can repeat a character in a new password. The attribute type is **SEC\_INT**.

**S\_MINAGE**

Defines the minimum age in weeks that the user's password must exist before the user can change it. The attribute type is **SEC\_INT**.

**S\_PWDCHECKS**

Defines the password restriction methods for this account. The attribute type is **SEC\_LIST**.

**S\_MINALPHA**

Defines the minimum number of alphabetic characters required in a new user's password. The attribute type is **SEC\_INT**.

**S\_MINDIFF**

Defines the minimum number of characters required in a new password that were not in the old password. The attribute type is **SEC\_INT**.

**S\_MINLEN**

Defines the minimum length of a user's password. The attribute type is **SEC\_INT**.

**S\_MINOTHER**

Defines the minimum number of non-alphabetic characters required in a new user's password. The attribute type is **SEC\_INT**.

**S\_DICTIIONLIST**

Defines the password dictionaries for this account. The attribute type is **SEC\_LIST**.

**S\_SUCHK**

Specifies whether the user account can be accessed with the **su** command. Type **SEC\_BOOL**.

**S\_REGISTRY**

Defines the user's authentication registry. The attribute type is **SEC\_CHAR**.



**S\_RLOGINCHK**

Specifies whether the user account can be used for remote logins using the **telnet** or **rlogin** commands. The attribute type is **SEC\_BOOL**.

**S\_DAEMONCHK**

Specifies whether the user account can be used for daemon execution of programs and subsystems using the **cron** daemon or **src**. The attribute type is **SEC\_BOOL**.

**S\_TPATH**

Defines how the account may be used on the trusted path. The attribute type is **SEC\_CHAR**. This attribute must be one of the following values:

**nosak** The secure attention key is not enabled for this account.

**notsh** The trusted shell cannot be accessed from this account.

**always**

This account may only run trusted programs.

**on** Normal trusted-path processing applies.

**S\_TTYS**

List of ttys that can or cannot be used to access this account. The attribute type is **SEC\_LIST**.

**S\_SUGROUPS**

Groups that can or cannot access this account. The attribute type is **SEC\_LIST**.

**S\_EXPIRATION**

Expiration date for this account is a string in the form MMDDhhmmyy, where MM is the month, DD is the day, hh is the hour in 0 to 24 hour notation, mm is the minutes past the hour, and yy is the last two digits of the year. The attribute type is **SEC\_CHAR**.

**S\_AUTH1**

Primary authentication methods for this account. The attribute type is **SEC\_LIST**.

**S\_AUTH2**

Secondary authentication methods for this account. The attribute type is **SEC\_LIST**.

**S\_UFSIZE**

Process file size soft limit. The attribute type is **SEC\_INT**.

**S\_UCPU**

Process CPU time soft limit. The attribute type is **SEC\_INT**.

**S\_UDATA**

Process data segment size soft limit. The attribute type is **SEC\_INT**.

**S\_USTACK**

Process stack segment size soft limit. Type: **SEC\_INT**.

**S\_URSS**

Process real memory size soft limit. Type: **SEC\_INT**.

**S\_UCORE**

Process core file size soft limit. The attribute type is **SEC\_INT**.

**S\_UNOFILE**

Process file descriptor table size soft limit. The attribute type is **SEC\_INT**.

**S\_PWD**

Specifies the value of the passwd field in the **/etc/passwd** file. The attribute type is **SEC\_CHAR**.

**S\_UMASK**

File creation mask for a user. The attribute type is **SEC\_INT**.

**S\_LOCKED**

Specifies whether the user's account can be logged into. The attribute type is **SEC\_BOOL**.

**S\_ROLES**

Defines the administrative roles for this account. The attribute type is **SEC\_LIST**.

**S\_UFSIZE\_HARD**

Process file size hard limit. The attribute type is **SEC\_INT**.

**S\_UCPU\_HARD**

Process CPU time hard limit. The attribute type is **SEC\_INT**.

**S\_UDATA\_HARD**

Process data segment size hard limit. The attribute type is **SEC\_INT**.

**S\_USREXPORT**

Specifies if the DCE registry can overwrite the local user information with the DCE user information during a DCE export operation. The attribute type is **SEC\_BOOL**.

**S\_USTACK\_HARD**

Process stack segment size hard limit. Type: **SEC\_INT**.

**S\_URSS\_HARD**

Process real memory size hard limit. Type: **SEC\_INT**.

**S\_UCORE\_HARD**

Process core file size hard limit. The attribute type is **SEC\_INT**.

**S\_UNOFILE\_HARD**

Process file descriptor table size hard limit. The attribute type is **SEC\_INT**.

**Note:** These values are string constants that should be used by applications both for convenience and to permit optimization in latter implementations. Additional user-defined attributes may be used and will be stored in the format specified by the *Type* parameter.

*Mode* Specifies the search mode. This parameter can be used to delimit the search to one or more user credentials databases. Specifying a non-null *Mode* value also implicitly rewinds the search. A null *Mode* value continues the search sequentially through the database. This parameter must include one of the following values specified as a bit mask; these are defined in the **usersec.h** file:

**S\_LOCAL**

Locally defined users are included in the search.

**S\_SYSTEM**

All credentials servers for the system are searched.

*Type* Specifies the type of attribute expected. Valid types are defined in the **usersec.h** file and include:

**SEC\_INT**

The format of the attribute is an integer.

For the **getuserattr** subroutine, the user should supply a pointer to a defined integer variable. For the **putuserattr** subroutine, the user should supply an integer.

**SEC\_CHAR**

The format of the attribute is a null-terminated character string.

For the **getuserattr** subroutine, the user should supply a pointer to a defined character pointer variable. For the **putuserattr** subroutine, the user should supply a character pointer.

### SEC\_LIST

The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series is terminated by two successive null characters.

For the **getuserattr** subroutine, the user should supply a pointer to a defined character pointer variable. For the **putuserattr** subroutine, the user should supply a character pointer.

### SEC\_BOOL

The format of the attribute from **getuserattr** is an integer with the value of either 0 (false) or 1 (true). The format of the attribute for **putuserattr** is a null-terminated string containing one of the following strings: true, false, yes, no, always, or never.

For the **getuserattr** subroutine, the user should supply a pointer to a defined integer variable. For the **putuserattr** subroutine, the user should supply a character pointer.

### SEC\_COMMIT

For the **putuserattr** subroutine, this value specified by itself indicates that changes to the named user are to be committed to permanent storage. The *Attribute* and *Value* parameters are ignored. If no user is specified, the changes to all modified users are committed to permanent storage.

### SEC\_DELETE

The corresponding attribute is deleted from the database.

### SEC\_NEW

Updates all the user database files with the new user name when using the **putuserattr** subroutine.

*UID* Specifies the user ID to be translated into a user name.

*User* Specifies the name of the user for which an attribute is to be read.

*Value* Specifies a buffer, a pointer to a buffer, or a pointer to a pointer depending on the *Attribute* and *Type* parameters. See the *Type* parameter for more details.

## Security

Files Accessed:

Mode	File
rw	/etc/passwd
rw	/etc/group
rw	/etc/security/user
rw	/etc/security/limits
rw	/etc/security/group
rw	/etc/security/environ

## Return Values

If successful, the **getuserattr** subroutine with the **S\_LOGINCHK** or **S\_RLOGINCHK** attribute specified and the **putuserattr** subroutine return 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error. For all other attributes, the **getuserattr** subroutine returns 0.

If successful, the **IDtouser** and **nextuser** subroutines return a character pointer to a buffer containing the requested user name. Otherwise, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

If any of these subroutines fail, the following is returned:

**EACCES** Access permission is denied for the data request.

If the **getuserattr** and **putuserattr** subroutines fail, one or more of the following is returned:

**ENOENT** The specified *User* parameter does not exist.  
**EINVAL** The *Attribute* parameter does not contain one of the defined attributes or null.  
**EINVAL** The *Value* parameter does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors may not be detected.  
**EPERM** Operation is not permitted.  
**ENOATTR** The specified attribute is not defined for this user.

If the **IDtouser** subroutine fails, one or more of the following is returned:

**ENOENT** The specified *User* parameter does not exist

If the **nextuser** subroutine fails, one or more of the following is returned:

**EINVAL** The *Mode* parameter is not one of null, **S\_LOCAL**, or **S\_SYSTEM**.  
**EINVAL** The *Argument* parameter is not null.  
**ENOENT** The end of the search was reached.

## Files

**/etc/passwd** Contains user IDs.

## Related Information

The “getgroupattr, IDtogroup, nextgroup, or putgroupattr Subroutine” on page 375, “getuserpw, putuserpw, or putuserpwhist Subroutine” on page 468, **setpwdb** subroutine, **setuserdb** subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getuserattr Subroutine

### Purpose

Retrieves multiple user attributes in the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int getuserattrs (User, Attributes, Count)
char * User;
dbattr_t * Attributes;
int Count
```

## Description

**Attention:** Do not use this subroutine and the **setpwent** and **setgrent** subroutines simultaneously. The results can be unpredictable.

The **getuserattrs** subroutine accesses user information. Because of its greater granularity and extensibility, use it instead of the **getpwent** routines.

The **getuserattrs** subroutine reads one or more attributes from the user database. If the database is not already open, this subroutine does an implicit open for reading. A call to the **getuserattrs** subroutine with an *Attributes* parameter of null and *Count* parameter of 0 for every new user verifies that the user exists.

The *Attributes* array contains information about each attribute that is to be read. The **dbattr\_t** data structure contains the following fields:

### **attr\_name**

The name of the desired attribute.

### **attr\_idx**

Used internally by the **getuserattrs** subroutine.

### **attr\_type**

The type of the desired attribute. The following possible attributes are defined in the **usersec.h** file:

#### **S\_CORECOMP**

Core compression status. The attribute type is **SEC\_CHAR**.

#### **S\_COREPATH**

Core path specification status. The attribute type is **SEC\_CHAR**.

#### **S\_COREPNAME**

Core path specification location. The attribute type is **SEC\_CHAR**.

#### **S\_CORENAMING**

Core naming status. The attribute type is **SEC\_CHAR**.

**S\_ID** User ID. The attribute type is **SEC\_INT**.

#### **S\_PGID**

Principle group ID. The attribute type is **SEC\_INT**.

#### **S\_PGRP**

Principle group name. The attribute type is **SEC\_CHAR**.

#### **S\_GROUPS**

Groups to which the user belongs. The attribute type is **SEC\_LIST**.

#### **S\_ADMGROUPS**

Groups for which the user is an administrator. The attribute type is **SEC\_LIST**.

#### **S\_ADMIN**

Administrative status of a user. The attribute type is **SEC\_BOOL**.

#### **S\_AUDITCLASSES**

Audit classes to which the user belongs. The attribute type is **SEC\_LIST**.

#### **S\_AUTHSYSTEM**

Defines the user's authentication method. The attribute type is **SEC\_CHAR**.

- S\_HOME**  
Home directory. The attribute type is **SEC\_CHAR**.
- S\_SHELL**  
Initial program run by a user. The attribute type is **SEC\_CHAR**.
- S\_GECOS**  
Personal information for a user. The attribute type is **SEC\_CHAR**.
- S\_USRENV**  
User-state environment variables. The attribute type is **SEC\_LIST**.
- S\_SYSENV**  
Protected-state environment variables. The attribute type is **SEC\_LIST**.
- S\_LOGINCHK**  
Specifies whether the user account can be used for local logins. The attribute type is **SEC\_BOOL**.
- S\_HISTEXPIRE**  
Defines the period of time (in weeks) that a user cannot reuse a password. The attribute type is **SEC\_INT**.
- S\_HISTSIZE**  
Specifies the number of previous passwords that the user cannot reuse. The attribute type is **SEC\_INT**.
- S\_MAXREPEAT**  
Defines the maximum number of times a user can repeat a character in a new password. The attribute type is **SEC\_INT**.
- S\_MINAGE**  
Defines the minimum age in weeks that the user's password must exist before the user can change it. The attribute type is **SEC\_INT**.
- S\_PWDCHECKS**  
Defines the password restriction methods for this account. The attribute type is **SEC\_LIST**.
- S\_MINALPHA**  
Defines the minimum number of alphabetic characters required in a new user's password. The attribute type is **SEC\_INT**.
- S\_MINDIFF**  
Defines the minimum number of characters required in a new password that were not in the old password. The attribute type is **SEC\_INT**.
- S\_MINLEN**  
Defines the minimum length of a user's password. The attribute type is **SEC\_INT**.
- S\_MINOTHER**  
Defines the minimum number of non-alphabetic characters required in a new user's password. The attribute type is **SEC\_INT**.
- S\_DICTIIONLIST**  
Defines the password dictionaries for this account. The attribute type is **SEC\_LIST**.
- S\_SUCHK**  
Specifies whether the user account can be accessed with the **su** command. Type **SEC\_BOOL**.
- S\_REGISTRY**  
Defines the user's authentication registry. The attribute type is **SEC\_CHAR**.

**S\_RLOGINCHK**

Specifies whether the user account can be used for remote logins using the **telnet** or **rlogin** commands. The attribute type is **SEC\_BOOL**.

**S\_DAEMONCHK**

Specifies whether the user account can be used for daemon execution of programs and subsystems using the **cron** daemon or **src**. The attribute type is **SEC\_BOOL**.

**S\_TPATH**

Defines how the account may be used on the trusted path. The attribute type is **SEC\_CHAR**. This attribute must be one of the following values:

**nosak** The secure attention key is not enabled for this account.

**notsh** The trusted shell cannot be accessed from this account.

**always**

This account may only run trusted programs.

**on** Normal trusted-path processing applies.

**S\_TTYS**

List of ttys that can or cannot be used to access this account. The attribute type is **SEC\_LIST**.

**S\_SUGROUPS**

Groups that can or cannot access this account. The attribute type is **SEC\_LIST**.

**S\_EXPIRATION**

Expiration date for this account is a string in the form MMDDhhmmyy, where MM is the month, DD is the day, hh is the hour in 0 to 24 hour notation, mm is the minutes past the hour, and yy is the last two digits of the year. The attribute type is **SEC\_CHAR**.

**S\_AUTH1**

Primary authentication methods for this account. The attribute type is **SEC\_LIST**.

**S\_AUTH2**

Secondary authentication methods for this account. The attribute type is **SEC\_LIST**.

**S\_UFSIZE**

Process file size soft limit. The attribute type is **SEC\_INT**.

**S\_UCPU**

Process CPU time soft limit. The attribute type is **SEC\_INT**.

**S\_UDATA**

Process data segment size soft limit. The attribute type is **SEC\_INT**.

**S\_USTACK**

Process stack segment size soft limit. Type: **SEC\_INT**.

**S\_URSS**

Process real memory size soft limit. Type: **SEC\_INT**.

**S\_UCORE**

Process core file size soft limit. The attribute type is **SEC\_INT**.

**S\_UNOFILE**

Process file descriptor table size soft limit. The attribute type is **SEC\_INT**.

**S\_PWD**

Specifies the value of the passwd field in the **/etc/passwd** file. The attribute type is **SEC\_CHAR**.

**S\_UMASK**

File creation mask for a user. The attribute type is **SEC\_INT**.

**S\_LOCKED**

Specifies whether the user's account can be logged into. The attribute type is **SEC\_BOOL**.

**S\_ROLES**

Defines the administrative roles for this account. The attribute type is **SEC\_LIST**.

**S\_UFSIZE\_HARD**

Process file size hard limit. The attribute type is **SEC\_INT**.

**S\_UCPU\_HARD**

Process CPU time hard limit. The attribute type is **SEC\_INT**.

**S\_UDATA\_HARD**

Process data segment size hard limit. The attribute type is **SEC\_INT**.

**S\_USREXPORT**

Specifies if the DCE registry can overwrite the local user information with the DCE user information during a DCE export operation. The attribute type is **SEC\_BOOL**.

**S\_USTACK\_HARD**

Process stack segment size hard limit. Type: **SEC\_INT**.

**S\_URSS\_HARD**

Process real memory size hard limit. Type: **SEC\_INT**.

**S\_UCORE\_HARD**

Process core file size hard limit. The attribute type is **SEC\_INT**.

**S\_UNOFILE\_HARD**

Process file descriptor table size hard limit. The attribute type is **SEC\_INT**.

**attr\_flag**

The results of the request to read the desired attribute.

**attr\_un**

A union containing the returned values. Its union members that follow correspond to the definitions of the **attr\_char**, **attr\_int**, **attr\_long**, and **attr\_llong** macros, respectively:

**un\_char**

Attributes of type **SEC\_CHAR** and **SEC\_LIST** store a pointer to the returned attribute in this member when the requested attribute is successfully read. The caller is responsible for freeing this memory.

**un\_int** Attributes of type **SEC\_INT** and **SEC\_BOOL** store the value of the attribute in this member when the requested attribute is successfully read.

**un\_long**

Attributes of type **SEC\_LONG** store the value of the attribute in this member when the requested attribute is successfully read.

**un\_llong**

Attributes of type **SEC\_LLONG** store the value of the attribute in this member when the requested attribute is successfully read.

**attr\_domain**

The authentication domain containing the attribute. The **getuserattns** subroutine is responsible for managing the memory referenced by this pointer.

If **attr\_domain** is specified for an attribute, the get request is sent only to that domain.

If **attr\_domain** is not specified (that is, set to NULL), **getuserattns** searches the domains known to the system and sets this field to the name of the domain from which the value is retrieved. This



search space can be restricted with the **setauthdb** subroutine so that only the domain specified in the **setauthdb** call is searched.

If the request for a NULL domain was not satisfied, the request is tried from the local files using the default stanza.

Use the **setuserdb** and **enduserdb** subroutines to open and close the user database. Failure to explicitly open and close the user database can result in loss of memory and performance.

## Parameters

<i>User</i>	Specifies the name of the user for which the attributes are to be read.
<i>Attributes</i>	A pointer to an array of zero or more elements of type <b>dbattr_t</b> . The list of user attributes is defined in the <b>usersec.h</b> header file.
<i>Count</i>	The number of array elements in <i>Attributes</i> .

## Security

Files accessed:

Mode	File
rw	<b>/etc/passwd</b>
rw	<b>/etc/group</b>
rw	<b>/etc/security/user</b>
rw	<b>/etc/security/limits</b>
rw	<b>/etc/security/group</b>
rw	<b>/etc/security/envIRON</b>

## Return Values

If *User* exists, the **getuserattrs** subroutine returns 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error. Each element in the *Attributes* array must be examined on a successful call to **getuserattrs** to determine if the *Attributes* array entry was successfully retrieved.

## Error Codes

The **getuserattrs** subroutine returns an error that indicates that the user does or does not exist. Additional errors can indicate an error querying the information databases for the requested attributes.

<b>EINVAL</b>	The <i>Count</i> parameter is less than 0.
<b>EINVAL</b>	The <i>Attributes</i> parameter is null and the <i>Count</i> parameter is greater than 0.
<b>ENOENT</b>	The specified <i>User</i> parameter does not exist.

If the **getuserattrs** subroutine fails to query an attribute, one or more of the following errors is returned in the **attr\_flag** field of the corresponding *Attributes* element:

<b>EACCES</b>	The user does not have access to the attribute specified in the <i>attr_name</i> field.
<b>EINVAL</b>	The <b>attr_type</b> field in the <i>Attributes</i> entry contains an invalid type.
<b>EINVAL</b>	The <b>attr_un</b> field in the <i>Attributes</i> entry does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors might not be detected.
<b>ENOATTR</b>	The <b>attr_name</b> field in the <i>Attributes</i> entry specifies an attribute that is not defined for this user or group.

## Examples

The following sample test program displays the output to a call to **getuserattrs**. In this example, the system has a user named foo.

```

#include <stdio.h>
#include <usersec.h>

#define NATTR 3
#define USERNAME "foo"

main() {

    dbattr_t attributes[NATTR];
    int i;
    int rc;

    memset (&attributes, 0, sizeof(attributes));

    /*
     * Fill in the attributes array with "id", "expires" and
     * "SYSTEM" attributes.
     */

    attributes[0].attr_name = S_ID;
    attributes[0].attr_type = SEC_INT;;

    attributes[1].attr_name = S_ADMIN;
    attributes[1].attr_type = SEC_BOOL;

    attributes[2].attr_name = S_AUTHSYSTEM;
    attributes[2].attr_type = SEC_CHAR;

    /*
     * Make a call to getuserattrs.
     */

        setuserdb(S_READ);

    rc = getuserattrs(USERNAME, attributes, NATTR);

        enduserdb();

    if (rc) {
        printf("getuserattrs failed ....\n");
        exit(-1);
    }

    for (i = 0; i < NATTR; i++) {
        printf("attribute name   : %s \n", attributes[i].attr_name);
        printf("attribute flag    : %d \n", attributes[i].attr_flag);

        if (attributes[i].attr_flag) {

            /*
             * No attribute value. Continue.
             */
            printf("\n");
            continue;
        }
        /*
         * We have a value.
         */
        printf("attribute domain : %s \n", attributes[i].attr_domain);
        printf("attribute value  : ");

        switch (attributes[i].attr_type)
        {
        case SEC_CHAR:
            if (attributes[i].attr_char) {
                printf("%s\n", attributes[i].attr_char);
            }
        }
    }
}

```

```

        free(attributes[i].attr_char);
    }
        break;
    case SEC_INT:
    case SEC_BOOL:
        printf("%d\n", attributes[i].attr_int);
        break;

    default:
        break;
    }
    printf("\n");
}
exit(0);
}

```

The following output for the call is expected:

```

attribute name   : id
attribute flag   : 0
attribute domain : files
attribute value  : 206

attribute name   : admin
attribute flag   : 0
attribute domain : files
attribute value  : 0

attribute name   : SYSTEM
attribute flag   : 0
attribute domain : files
attribute value  : compat

```

## Files

**/etc/passwd**                      Contains user IDs.

## Related Information

The “getgroupattrs Subroutine” on page 378, “getuserpw, putuserpw, or putuserpwhist Subroutine” on page 468, setpwdb Subroutine, setuserdb Subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## GetUserAuths Subroutine

### Purpose

Accesses the set of authorizations of a user.

### Library

Security Library (**libc.a**)

### Syntax

```

#include <usersec.h>
char *GetUserAuths(void);

```

## Description

The **GetUserAuths** subroutine returns the list of authorizations associated with the real user ID and group set of the process. By default, the ALL authorization is returned for the root user.

## Return Values

If successful, the **GetUserAuths** subroutine returns a list of authorizations associated with the user. The format of the list is a series of concatenated strings, each null-terminated. A null string terminates the list. Otherwise, a null pointer is returned and the **errno** global variable is set to indicate the error.

---

## getuserpw, putuserpw, or putuserpwhist Subroutine

### Purpose

Accesses the user authentication data.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <userpw.h>
```

```
struct userpw *getuserpw ( User)
char *User;
```

```
int putuserpw ( Password)
struct userpw *Password;
```

```
int putuserpwhist ( Password, Message)
struct userpw *Password;
char **Message;
```

## Description

These subroutines may be used to access user authentication information. Because of their greater granularity and extensibility, you should use them instead of the **getpwent** routines.

The **getuserpw** subroutine reads the user's locally defined password information. If the **setpwdb** subroutine has not been called, the **getuserpw** subroutine will call it as **setpwdb** (**S\_READ**). This can cause problems if the **putuserpw** subroutine is called later in the program.

The **putuserpw** subroutine updates or creates a locally defined password information stanza in the **/etc/security/passwd** file. The password entry created by the **putuserpw** subroutine is used only if there is an ! (exclamation point) in the **/etc/passwd** file's password field. The user application can use the **putuserattr** subroutine to add an ! to this field.

The **putuserpw** subroutine will open the authentication database read/write if no other access has taken place, but the program should call **setpwdb** (**S\_READ** | **S\_WRITE**) before calling the **putuserpw** subroutine.

The **putuserpwhist** subroutine updates or creates a locally defined password information stanza in the **etc/security/passwd** file. The subroutine also manages a database of previous passwords used for password reuse restriction checking. It is recommended to use the **putuserpwhist** subroutine, rather than the **putuserpw** subroutine, to ensure the password is added to the password history database.

## Parameters

<i>Password</i>	<p>Specifies the password structure used to update the password information for this user. This structure is defined in the <b>userpw.h</b> file and contains the following members:</p> <p><b>upw_name</b> Specifies the user's name. (The first eight characters must be unique, since longer names are truncated.)</p> <p><b>upw_passwd</b> Specifies the user's password.</p> <p><b>upw_lastupdate</b> Specifies the time, in seconds, since the epoch (that is, 00:00:00 GMT, January 1, 1970), when the password was last updated.</p> <p><b>upw_flags</b> Specifies attributes of the password. This member is a bit mask of one or more of the following values, defined in the <b>userpw.h</b> file.</p> <p><b>PW_NOCHECK</b> Specifies that new passwords need not meet password restrictions in effect for the system.</p> <p><b>PW_ADMCHG</b> Specifies that the password was last set by an administrator and must be changed at the next successful use of the <b>login</b> or <b>su</b> command.</p> <p><b>PW_ADMIN</b> Specifies that password information for this user may only be changed by the root user.</p>
<i>Message</i>	<p>Indicates a message that specifies an error occurred while updating the password history database. Upon return, the value is either a pointer to a valid string within the memory allocated storage or a null pointer.</p>
<i>User</i>	<p>Specifies the name of the user for which password information is read. (The first eight characters must be unique, since longer names are truncated.)</p>

## Security

Files Accessed:

Mode	File
rw	<b>/etc/security/passwd</b>

## Return Values

If successful, the **getuserpw** subroutine returns a valid pointer to a **userpw** structure. Otherwise, a null pointer is returned and the **errno** global variable is set to indicate the error. If the user exists but there is no user entry in the **/etc/security/passwd** file, the **getuserpw** subroutine returns success with the name field set to user name, the password field set to NULL, the lastupdate field set to 0 and the flags field set to 0. If the user exists and there is an entry in the **/etc/security/passwd** file but one or more fields are missing, the **getuserpw** subroutine returns the fields that exist.

If the user exists but there is no user entry in the **/etc/security/passwd** file, the **putuserpw** subroutine creates a user stanza in the **/etc/security/passwd** file. If the user exists and there is an entry in the **/etc/security/passwd** file but one or more fields are missing, the **putuserpw** subroutine updates the fields that exist and creates the fields that are missing.

If successful, the **putuserpwhist** subroutine returns a value of 0. If the subroutine failed to update or create a locally defined password information stanza in the **/etc/security/passwd** file, the **putuserpwhist** subroutine returns a nonzero value. If the subroutine was unable to update the password history database, a message is returned in the *Message* parameter and a return code of 0 is returned. If the user exists but there is no user entry in the **/etc/security/passwd** file, the **putuserpwhist** subroutine creates a user stanza in the **/etc/security/passwd** file and updates the password history. If the user exists and there is an entry in the **/etc/security/passwd** file but one or more fields are missing, the **putuserpwhist** subroutine updates the fields that exist, creates the fields that are missing and modifies the password history.

## Error Codes

The **getuserpw**, **putuserpw**, and **putuserpwhist** subroutines fail if the following values are true:

<b>EACCES</b>	The user is not able to open the files that contain the password attributes.
<b>ENOENT</b>	The user does not exist in the <b>/etc/passwd</b> file.

Subroutines invoked by the **getuserpw**, **putuserpw**, or **putuserpwhist** subroutines can also set errors.

## Files

<b>/etc/security/passwd</b>	Contains user passwords.
-----------------------------	--------------------------

## Related Information

The “getgroupattr, IDtogroup, nextgroup, or putgroupattr Subroutine” on page 375, “getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454, **setpwdb** or **endpwdb** subroutine, **setuserdb** subroutine.

List of Security and Auditing Subroutines and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getuserpwx Subroutine

### Purpose

Accesses the user authentication data.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <userpw.h>
```

```
struct userpwx *getuserpwx (User)
char * User;
```

### Description

The **getuserpwx** subroutine accesses user authentication information. Because of its greater granularity and extensibility, use it instead of the **getpwent** routines.

The **getuserpwx** subroutine reads the user’s password information from the local administrative domain or from a loadable authentication module that supports the required user attributes.

The **getuserpw** subroutine opens the authentication database read-only if no other access has taken place, but the program should call **setpwdb (S\_READ)** followed by **endpwdb** after access to the authentication database is no longer required.

The data returned by **getuserpwx** is stored in allocated memory and must be freed by the caller when the data is no longer required. The entire structure can be freed by invoking the **free** subroutine with the pointer returned by **getuserpwx**.

## Parameters

*User* Specifies the name of the user for which password information is read.

## Security

Files accessed:

Mode	File
r	/etc/passwd
r	/etc/security/passwd

## Return Values

If successful, the **getuserpwx** subroutine returns a valid pointer to a **userpwx** structure. Otherwise, a null pointer is returned and the **errno** global variable is set to indicate the error. The fields in a **userpwx** structure are defined in the **userpw.h** file, and they include the following members:

<b>upw_name</b>	Specifies the user's name.
<b>upw_passwd</b>	Specifies the user's encrypted password.
<b>upw_lastupdate</b>	Specifies the time, in seconds, since the epoch (that is, 00:00:00 GMT, 1 January 1970), when the password was last updated.
<b>upw_flags</b>	Specifies attributes of the password. This member is a bit mask of one or more of the following values, defined in the <b>userpw.h</b> file:  <b>PW_NOCHECK</b> Specifies that new passwords need not meet password restrictions in effect for the system.  <b>PW_ADMCHG</b> Specifies that the password was last set by an administrator and must be changed at the next successful use of the login or <b>su</b> command.  <b>PW_ADMIN</b> Specifies that password information for this user can only be changed by the root user.
<b>upw_authdb</b>	Specifies the administrative domain containing the authentication data.

## Error Codes

The **getuserpwx** subroutine fails if one of the following values is true:

<b>EACCES</b>	The user is not able to open the files that contain the password attributes.
<b>ENOENT</b>	The user does not have an entry in the <b>/etc/security/passwd</b> file or other administrative domain.

Subroutines invoked by the **getuserpwx** subroutine can also set errors.

## Files

`/etc/security/passwd`      Contains user passwords.

## Related Information

The “getgroupattrs Subroutine” on page 378, “getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454, “getuserattr Subroutine” on page 460, setpwdb Subroutinesetuserdb Subroutine.

---

## getusraclattr, nextusracl or putusraclattr Subroutine

### Purpose

Accesses the user screen information in the SMIT ACL database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>

int getusraclattr(User, Attribute, Value, Type)
char *User;
char *Attribute;
void *Value;
int Type;

char *nextusracl(void)

int putusraclattr(User, Attribute, Value, Type)
char *User;
char *Attribute;
void *Value;
int Type;
```

### Description

The **getusraclattr** subroutine reads a specified user attribute from the SMIT ACL database. If the database is not already open, this subroutine does an implicit open for reading.

Similarly, the **putusraclattr** subroutine writes a specified attribute into the user SMIT ACL database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by the **putusraclattr** subroutine must be explicitly committed by calling the **putusraclattr** subroutine with a *Type* parameter specifying **SEC\_COMMIT**. Until all the data is committed, only the **getusraclattr** subroutine within the process returns written data.

The **nextusracl** subroutine returns the next user in a linear search of the user SMIT ACL database. The consistency of consecutive searches depends upon the underlying storage-access mechanism and is not guaranteed by this subroutine.

The **setacldb** and **endacldb** subroutines should be used to open and close the database.



## Parameters

<i>Attribute</i>	Specifies which attribute is read. The following possible attributes are defined in the <b>usersec.h</b> file: <b>S_SCREEN</b> String of SMIT screens. The attribute type is <b>SEC_LIST</b> . <b>S_ACLMODE</b> String specifying the SMIT ACL database search scope. The attribute type is <b>SEC_CHAR</b> . <b>S_FUNCMODE</b> String specifying the databases to be searched. The attribute type is <b>SEC_CHAR</b> .
<i>Type</i>	Specifies the type of attribute expected. Valid types are defined in the <b>usersec.h</b> file and include: <b>SEC_CHAR</b> The format of the attribute is a null-terminated character string. For the <b>getusraclattr</b> subroutine, the user should supply a pointer to a defined character pointer variable. For the <b>putusraclattr</b> subroutine, the user should supply a character pointer. <b>SEC_LIST</b> The format of the attribute is a series of concatenated strings, each null-terminated. The last string in the series must be an empty (zero character count) string. For the <b>getusraclattr</b> subroutine, the user should supply a pointer to a defined character pointer variable. For the <b>putusraclattr</b> subroutine, the user should supply a character pointer. <b>SEC_COMMIT</b> For the <b>putusraclattr</b> subroutine, this value specified by itself indicates that changes to the named user are to be committed to permanent storage. The <i>Attribute</i> and <i>Value</i> parameters are ignored. If no user is specified, the changes to all modified users are committed to permanent storage. <b>SEC_DELETE</b> The corresponding attribute is deleted from the user SMIT ACL database. <b>SEC_NEW</b> Updates the user SMIT ACL database file with the new user name when using the <b>putusraclattr</b> subroutine.
<i>Value</i>	Specifies a buffer, a pointer to a buffer, or a pointer to a pointer depending on the <i>Attribute</i> and <i>Type</i> parameters. See the <i>Type</i> parameter for more details.

## Return Values

If successful, the **getusraclattr** returns 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

Possible return codes are:

<b>EACCES</b>	Access permission is denied for the data request.
<b>ENOENT</b>	The specified User parameter does not exist or the attribute is not defined for this user.
<b>ENOATTR</b>	The specified user attribute does not exist for this user.
<b>EINVAL</b>	The <i>Attribute</i> parameter does not contain one of the defined attributes or null.
<b>EINVAL</b>	The <i>Value</i> parameter does not point to a valid buffer or to valid data for this type of attribute.
<b>EPERM</b>	Operation is not permitted.

## Related Information

The **getgrpaclattr**, **nextgrpacl**, or **putgrpaclattr** (“getgrpaclattr, nextgrpacl, or putgrpaclattr Subroutine” on page 384) subroutine, **setacldb**, or **endacldb** subroutine.

---

## getutent, getutid, getutline, pututline, setutent, endutent, or utmpname Subroutine

### Purpose

Accesses **utmp** file entries.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <utmp.h>
```

```
struct utmp *getutent ( )
```

```
struct utmp *getutid ( ID)
struct utmp *ID;
```

```
struct utmp *getutline ( Line)
struct utmp *Line;
```

```
void pututline ( Utmp)
struct utmp *Utmp;
```

```
void setutent ( )
```

```
void endutent ( )
```

```
void utmpname ( File)
char *File;
```

### Description

The **getutent**, **getutid**, and **getutline** subroutines return a pointer to a structure of the following type:

```
struct utmp
> {
>     char ut_user[256];           /* User name */
>     char ut_id[14];             /* /etc/inittab ID */
>     char ut_line[64];          /* Device name (console, lnx) */
>     pid_t ut_pid;              /* Process ID */
>     short ut_type;             /* Type of entry */
>     int __time_t_space;        /* for 32vs64-bit time_t PPC */
>     time_t ut_time;           /* time entry was made */
>     struct exit_status
>     {
>         short e_termination; /* Process termination status */
>         short e_exit;        /* Process exit status */
>     }
>     ut_exit;                   /* The exit status of a process
>     /* marked as DEAD_PROCESS. */
>     char ut_host[256];         /* host name */
```

```

>     int __dbl_word_pad;      /* for double word alignment */
>     int __reservedA[2];
>     int __reservedV[6];
> };

```

The **getutent** subroutine reads the next entry from a **utmp**-like file. If the file is not open, this subroutine opens it. If the end of the file is reached, the **getutent** subroutine fails.

The **pututline** subroutine writes the supplied *Utmp* parameter structure into the **utmp** file. It is assumed that the user of the **pututline** subroutine has searched for the proper entry point using one of the **getut** subroutines. If not, the **pututline** subroutine calls **getutid** to search forward for the proper place. If so, **pututline** does not search. If the **pututline** subroutine does not find a matching slot for the entry, it adds a new entry to the end of the file.

The **setutent** subroutine resets the input stream to the beginning of the file. Issue a **setuid** call before each search for a new entry if you want to examine the entire file.

The **endutent** subroutine closes the file currently open.

The **utmpname** subroutine changes the name of a file to be examined from **/etc/utmp** to any other file. The name specified is usually **/var/adm/wtmp**. If the specified file does not exist, no indication is given. You are not aware of this fact until your first attempt to reference the file. The **utmpname** subroutine does not open the file. It closes the old file, if currently open, and saves the new file name.

The most current entry is saved in a static structure. To make multiple accesses, you must copy or use the structure between each access. The **getutid** and **getutline** subroutines examine the static structure first. If the contents of the static structure match what they are searching for, they do not read the **utmp** file. Therefore, you must fill the static structure with zeros after each use if you want to use these subroutines to search for multiple occurrences.

If the **pututline** subroutine finds that it is not already at the correct place in the file, the implicit read it performs does not overwrite the contents of the static structure returned by the **getutent** subroutine, the **getutid** subroutine, or the **getutline** subroutine. This allows you to get an entry with one of these subroutines, modify the structure, and pass the pointer back to the **pututline** subroutine for writing.

These subroutines use buffered standard I/O for input. However, the **pututline** subroutine uses an unbuffered nonstandard write to avoid race conditions between processes trying to modify the **utmp** and **wtmp** files.

## Parameters

<i>ID</i>	If you specify a type of <b>RUN_LVL</b> , <b>BOOT_TIME</b> , <b>OLD_TIME</b> , or <b>NEW_TIME</b> in the <i>ID</i> parameter, the <b>getutid</b> subroutine searches forward from the current point in the <b>utmp</b> file until an entry with a <i>ut_type</i> matching <i>ID</i> -> <i>ut_type</i> is found. If you specify a type of <b>INIT_PROCESS</b> , <b>LOGIN_PROCESS</b> , <b>USER_PROCESS</b> , or <b>DEAD_PROCESS</b> in the <i>ID</i> parameter, the <b>getutid</b> subroutine returns a pointer to the first entry whose type is one of these four and whose <i>ut_id</i> field matches <i>Id</i> -> <i>ut_id</i> . If the end of the file is reached without a match, the <b>getutid</b> subroutine fails.
<i>Line</i>	The <b>getutline</b> subroutine searches forward from the current point in the <b>utmp</b> file until it finds an entry of type <b>LOGIN_PROCESS</b> or <b>USER_PROCESS</b> that also has a <i>ut_line</i> string matching the <i>Line</i> -> <i>ut_line</i> parameter string. If the end of file is reached without a match, the <b>getutline</b> subroutine fails.
<i>Utmp</i>	Points to the <b>utmp</b> structure.
<i>File</i>	Specifies the name of the file to be examined.

## Return Values

These subroutines fail and return a null pointer if a read or write fails due to a permission conflict or because the end of the file is reached.

## Files

<code>/etc/utmp</code>	Path to the <b>utmp</b> file, which contains a record of users logged into the system.
<code>/var/adm/wtmp</code>	Path to the <b>wtmp</b> file, which contains accounting information about users logged in.

## Related Information

The **ttyslot** subroutine.

The **failedlogin**, **utmp**, or **wtmp** file.

---

## getvfsent, getvfsbytype, getvfsbyname, getvfsbyflag, setvfsent, or endvfsent Subroutine

### Purpose

Gets a **vfs** file entry.

### Library

Standard C Library(**libc.a**)

### Syntax

```
#include <sys/vfs.h>
#include <sys/vmount.h>
struct vfs_ent *getvfsent( )

struct vfs_ent *getvfsbytype( vfsType)
int vfsType;

struct vfs_ent *getvfsbyname( vfsName)
char *vfsName;

struct vfs_ent *getvfsbyflag( vfsFlag)
int vfsFlag;
void setvfsent( )
void endvfsent( )
```

### Description

**Attention:** All information is contained in a static area and so must be copied to be saved.

The **getvfsent** subroutine, when first called, returns a pointer to the first **vfs\_ent** structure in the file. On the next call, it returns a pointer to the next **vfs\_ent** structure in the file. Successive calls are used to search the entire file.

The **vfs\_ent** structure is defined in the **vfs.h** file and it contains the following fields:

```

char vfsent_name;
int vfsent_type;
int vfsent_flags;
char *vfsent_mnt_hlpr;
char *vfsent_fs_hlpr;

```

The **getvfsbytype** subroutine searches from the beginning of the file until it finds a **vfs** type matching the *vfsType* parameter. The subroutine then returns a pointer to the structure in which it was found.

The **getvfsbyname** subroutine searches from the beginning of the file until it finds a **vfs** name matching the *vfsName* parameter. The search is made using flattened names; the search-string uses ASCII equivalent characters.

The **getvfsbytype** subroutine searches from the beginning of the file until it finds a type matching the *vfsType* parameter.

The **getvfsbyflag** subroutine searches from the beginning of the file until it finds the entry whose flag corresponds flags defined in the **vfs.h** file. Currently, these are **VFS\_DFLT\_LOCAL** and **VFS\_DFLT\_REMOTE**.

The **setvfsent** subroutine rewinds the **vfs** file to allow repeated searches.

The **endvfsent** subroutine closes the **vfs** file when processing is complete.

## Parameters

<i>vfsType</i>	Specifies a <b>vfs</b> type.
<i>vfsName</i>	Specifies a <b>vfs</b> name.
<i>vfsFlag</i>	Specifies either <b>VFS_DFLT_LOCAL</b> or <b>VFS_DFLT_REMOTE</b> .

## Return Values

The **getvfsent**, **getvfsbytype**, **getvfsbyname**, and **getvfsbyflag** subroutines return a pointer to a **vfs\_ent** structure containing the broken-out fields of a line in the **/etc/vfs** file. If an end-of-file character or an error is encountered on reading, a null pointer is returned.

## Files

<b>/etc/vfs</b>	Describes the virtual file system (VFS) installed on the system.
-----------------	--

## Related Information

The **getvfsent**, **getvfsbytype**, **getvfsbyname**, **getvfsbyflag**, **setvfsent**, or **endvfsent** (“getvfsent, getvfsbytype, getvfsbyname, getvfsbyflag, setvfsent, or endvfsent Subroutine” on page 369) subroutine.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## getwc, fgetwc, or getwchar Subroutine

### Purpose

Gets a wide character from an input stream.

## Library

Standard I/O Package (**libc.a**)

## Syntax

```
#include <stdio.h>

wint_t getwc ( Stream)
FILE *Stream;

wint_t fgetwc (Stream)
FILE *Stream;

wint_t getwchar (void)
```

## Description

The **fgetwc** subroutine obtains the next wide character from the input stream specified by the *Stream* parameter, converts it to the corresponding wide character code, and advances the file position indicator the number of bytes corresponding to the obtained multibyte character. The **getwc** subroutine is equivalent to the **fgetwc** subroutine, except that when implemented as a macro, it may evaluate the *Stream* parameter more than once. The **getwchar** subroutine is equivalent to the **getwc** subroutine with **stdin** (the standard input stream).

The first successful run of the **fgetc** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **fgets** (“gets or fgets Subroutine” on page 434), **fgetwc**, **fgetws** (“getws or fgets Subroutine” on page 480), **fread** (“fread or fwrite Subroutine” on page 310), **fscanf**, **getc** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **getchar** (“getc, getchar, fgetc, or getw Subroutine” on page 347), **gets** (“gets or fgets Subroutine” on page 434), or **scanf** subroutine using a stream that returns data not supplied by a prior call to the **ungetc** or **ungetwc** subroutine marks the *st\_atime* field for update.

## Parameters

*Stream*                      Specifies input data.

## Return Values

Upon successful completion, the **getwc** and **fgetwc** subroutines return the next wide character from the input stream pointed to by the *Stream* parameter. The **getwchar** subroutine returns the next wide character from the input stream pointed to by **stdin**.

If the end of the file is reached, an indicator is set and **WEOF** is returned. If a read error occurs, an error indicator is set, **WEOF** is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

If the **getwc**, **fgetwc**, or **getwchar** subroutine is unsuccessful because the stream is not buffered or data needs to be read into the buffer, it returns one of the following error codes:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> flag is set for the file descriptor underlying the <i>Stream</i> parameter, delaying the process.
<b>EBADF</b>	Indicates that the file descriptor underlying the <i>Stream</i> parameter is not valid and cannot be opened for reading.
<b>EINTR</b>	Indicates that the process has received a signal that terminates the read operation.
<b>EIO</b>	Indicates that a physical error has occurred, or the process is in a background process group attempting to read from the controlling terminal, and either the process is ignoring or blocking the <b>SIGTTIN</b> signal or the process group is orphaned.
<b>EOVERFLOW</b>	Indicates that the file is a regular file and an attempt has been made to read at or beyond the offset maximum associated with the corresponding stream.

The **getwc**, **fgetwc**, or **getwchar** subroutine is also unsuccessful due to the following error conditions:

<b>ENOMEM</b>	Indicates that storage space is insufficient.
<b>ENXIO</b>	Indicates that the process sent a request to a nonexistent device, or the device cannot handle the request.
<b>EILSEQ</b>	Indicates that the <b>wc</b> wide-character code does not correspond to a valid character.

## Related Information

Other wide character I/O subroutines: **getws** or **fgetws** (“getws or fgetws Subroutine” on page 480) subroutine, **putwc**, **putwchar**, or **fputwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **putws** or **fputws** (“putws or fputws Subroutine” on page 1341) subroutine, **ungetwc** subroutine.

Related standard I/O subroutines: **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **gets** or **fgets** (“gets or fgets Subroutine” on page 434) subroutine, **fread** (“fread or fwrite Subroutine” on page 310) subroutine, **fwrite** (“fread or fwrite Subroutine” on page 310) subroutine, **printf**, **fprintf**, **sprintf**, **wsprintf**, **vprintf**, **vfprintf**, **vsprintf**, or **vwsprintf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **putc**, **putchar**, **fputc**, or **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **puts** or **fputs** (“puts or fputs Subroutine” on page 1331) subroutine.

Subroutines, Example Programs, and Libraries and Understanding Wide Character Input/Output Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## getwd Subroutine

### Purpose

Gets current directory path name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
char *getwd ( PathName)
```

```
char *PathName;
```

### Description

The **getwd** subroutine determines the absolute path name of the current directory, then copies that path name into the area pointed to by the *PathName* parameter.

The maximum path-name length, in characters, is set by the **PATH\_MAX** value, as specified in the **limits.h** file.

## Parameters

*PathName* Points to the full path name.

## Return Values

If the call to the **getwd** subroutine is successful, a pointer to the absolute path name of the current directory is returned. If an error occurs, the **getwd** subroutine returns a null value and places an error message in the *PathName* parameter.

In UNIX03 mode, the **getwd** subroutine returns a null value if the actual path name is longer than the value defined by **PATH\_MAX**. In the previous mode, the **getwd** subroutine returns a truncated path name if the path name is longer than **PATH\_MAX**. The previous behavior can be disabled setting the environment variable **XPG\_SUS\_ENV=ON**.

## Related Information

The “getcwd Subroutine” on page 358.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## getws or fgetws Subroutine

### Purpose

Gets a string from a stream.

### Library

Standard I/O Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
wchar_t *fgetws ( WString, Number, Stream)
```

```
wchar_t *WString;
```

```
int Number;
```

```
FILE *Stream;
```

```
wchar_t *getws (WString)
```

```
wchar_t *WString;
```

### Description

The **fgetws** subroutine reads characters from the input stream, converts them to the corresponding wide character codes, and places them in the array pointed to by the *WString* parameter. The subroutine continues until either the number of characters specified by the *Number* parameter minus 1 are read or the subroutine encounters a new-line or end-of-file character. The **fgetws** subroutine terminates the wide character string specified by the *WString* parameter with a null wide character.

The **getws** subroutine reads wide characters from the input stream pointed to by the standard input stream (**stdin**) into the array pointed to by the *WString* parameter. The subroutine continues until it encounters a new-line or the end-of-file character, then it discards any new-line character and places a null wide character after the last character read into the array.



## Parameters

<i>WString</i>	Points to a string to receive characters.
<i>Stream</i>	Points to the <b>FILE</b> structure of an open file.
<i>Number</i>	Specifies the maximum number of characters to read.

## Return Values

If the **getws** or **fgetws** subroutine reaches the end of the file without reading any characters, it transfers no characters to the *String* parameter and returns a null pointer. If a read error occurs, the **getws** or **fgetws** subroutine returns a null pointer and sets the **errno** global variable to indicate the error.

## Error Codes

If the **getws** or **fgetws** subroutine is unsuccessful because the stream is not buffered or data needs to be read into the stream's buffer, it returns one or more of the following error codes:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> flag is set for the file descriptor underlying the <i>Stream</i> parameter, and the process is delayed in the <b>fgetws</b> subroutine.
<b>EBADF</b>	Indicates that the file descriptor specifying the <i>Stream</i> parameter is not a read-access file.
<b>EINTR</b>	Indicates that the read operation is terminated due to the receipt of a signal, and either no data was transferred or the implementation does not report partial transfer for this file.
<b>EIO</b>	Indicates that insufficient storage space is available.
<b>ENOMEM</b>	Indicates that insufficient storage space is available.
<b>EILSEQ</b>	Indicates that the data read from the input stream does not form a valid character.

## Related Information

Other wide character I/O subroutines: **fgetwc** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **fputwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **fputws** (“putws or fputws Subroutine” on page 1341) subroutine, **getwc** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **putwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **putwchar** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **putws** (“putws or fputws Subroutine” on page 1341) subroutine, **ungetwc** subroutine.

Related standard I/O subroutines: **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fgetc** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **fgets** (“gets or fgets Subroutine” on page 434) subroutine, **fopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fprintf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **fputc** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **fputs** (“puts or fputs Subroutine” on page 1331) subroutine, **fread** (“fread or fwrite Subroutine” on page 310) subroutine, **freopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fscanf** subroutine, **fwrite** (“fread or fwrite Subroutine” on page 310) subroutine, **getc** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **getchar** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **gets** (“gets or fgets Subroutine” on page 434) subroutine, **printf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **putc** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **putchar** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **puts** (“puts or fputs Subroutine” on page 1331) subroutine, **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **scanf** subroutine, **sprintf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **ungetc** subroutine.

Understanding Wide Character Input/Output Subroutines and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

## glob Subroutine

### Purpose

Generates path names.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <glob.h>
```

```
int glob (Pattern, Flags, (Errfunc)(), Pglob)
const char *Pattern;
int Flags;
int *Errfunc (Epath, Eerrno)
const char *Epath;
int Eerrno;
glob_t *Pglob;
```

### Description

The **glob** subroutine constructs a list of accessible files that match the *Pattern* parameter.

The **glob** subroutine matches all accessible path names against this pattern and develops a list of all matching path names. To have access to a path name, the **glob** subroutine requires search permission on every component of a path except the last, and read permission on each directory of any file name component of the *Pattern* parameter that contains any of the special characters \* (asterisk), ? (question mark), or [ (left bracket). The **glob** subroutine stores the number of matched path names and a pointer to a list of pointers to path names in the *Pglob* parameter. The path names are in sort order, based on the setting of the **LC\_COLLATE** category in the current locale. The first pointer after the last path name is a null character. If the pattern does not match any path names, the returned number of matched paths is zero.

### Parameters

#### *Pattern*

Contains the file name pattern to compare against accessible path names.

*Flags* Controls the customizable behavior of the **glob** subroutine.

The *Flags* parameter controls the behavior of the **glob** subroutine. The *Flags* value is the bitwise inclusive OR of any of the following constants, which are defined in the **glob.h** file:

#### **GLOB\_APPEND**

Appends path names located with this call to any path names previously located. If the **GLOB\_APPEND** constant is not set, new path names overwrite previous entries in the *Pglob* array. The **GLOB\_APPEND** constant should not be set on the first call to the **glob** subroutine. It may, however, be set on subsequent calls.

The **GLOB\_APPEND** flag can be used to append a new set of path names to those found in a previous call to the **glob** subroutine. If the **GLOB\_APPEND** flag is specified in the *Flags* parameter, the following rules apply:

- If the application sets the **GLOB\_DOOFFS** flag in the first call to the **glob** subroutine, it is also set in the second. The value of the *Pglob* parameter is not modified between the calls.
- If the application did not set the **GLOB\_DOOFFS** flag in the first call to the **glob** subroutine, it is not set in the second.
- After the second call, the *Pglob* parameter points to a list containing the following:
  - Zero or more null characters, as specified by the **GLOB\_DOOFFS** flag.
  - Pointers to the path names that were in the *Pglob* list before the call, in the same order as after the first call to the **glob** subroutine.
  - Pointers to the new path names generated by the second call, in the specified order.
- The count returned in the *Pglob* parameter is the total number of path names from the two calls.
- The application should not modify the *Pglob* parameter between the two calls.

It is the caller's responsibility to create the structure pointed to by the *Pglob* parameter. The **glob** subroutine allocates other space as needed.

#### **GLOB\_DOOFFS**

Uses the **gl\_offs** structure to specify the number of null pointers to add to the beginning of the **gl\_pathv** component of the *Pglob* parameter.

#### **GLOB\_ERR**

Causes the **glob** subroutine to return when it encounters a directory that it cannot open or read. If the **GLOB\_ERR** flag is not set, the **glob** subroutine continues to find matches if it encounters a directory that it cannot open or read.

#### **GLOB\_MARK**

Specifies that each path name that is a directory should have a / (slash) appended.

#### **GLOB\_NOCHECK**

If the *Pattern* parameter does not match any path name, the **glob** subroutine returns a list consisting only of the *Pattern* parameter, and the number of matched patterns is one.

#### **GLOB\_NOSORT**

Specifies that the list of path names need not be sorted. If the **GLOB\_NOSORT** flag is not set, path names are collated according to the current locale.

#### **GLOB\_QUOTE**

If the **GLOB\_QUOTE** flag is set, a \ (backslash) can be used to escape metacharacters.

#### *Errfunc*

Specifies an optional subroutine that, if specified, is called when the **glob** subroutine detects an error condition.

#### *Pglob*

Contains a pointer to a **glob\_t** structure. The structure is allocated by the caller. The array of structures containing the file names matching the *Pattern* parameter are defined by the **glob** subroutine. The last entry is a null pointer.

#### *Epath*

Specifies the path that failed because a directory could not be opened or read.

#### *Eerrno*

Specifies the **errno** value of the failure indicated by the *Epath* parameter. This value is set by the **opendir**, **readdir**, or **stat** subroutines.

## Return Values

On successful completion, the **glob** subroutine returns a value of 0. The *Pglob* parameter returns the number of matched path names and a pointer to a null-terminated list of matched and sorted path names. If the number of matched path names in the *Pglob* parameter is zero, the pointer in the *Pglob* parameter is undefined.

## Error Codes

If the **glob** subroutine terminates due to an error, it returns one of the nonzero constants below. These are defined in the **glob.h** file. In this case, the *Pglob* values are still set as defined in the Return Values section.

<b>GLOB_ABORTED</b>	Indicates the scan was stopped because the <b>GLOB_ERROR</b> flag was set or the subroutine specified by the <b>errfunc</b> parameter returned a nonzero value.
<b>GLOB_NOSPACE</b>	Indicates a failed attempt to allocate memory.

If, during the search, a directory is encountered that cannot be opened or read and the *Errfunc* parameter is not a null value, the **glob** subroutine calls the subroutine specified by the **errfunc** parameter with two arguments:

- The *Epath* parameter specifies the path that failed.
- The *Eerrno* parameter specifies the value of the **errno** global variable from the failure, as set by the **opendir**, **readdir**, or **stat** subroutine.

If the subroutine specified by the *Errfunc* parameter is called and returns nonzero, or if the **GLOB\_ERR** flag is set in the *Flags* parameter, the **glob** subroutine stops the scan and returns **GLOB\_ABORTED** after setting the *Pglob* parameter to reflect the paths already scanned. If **GLOB\_ERR** is not set and either the *Errfunc* parameter is null or *\*errfunc* returns zero, the error is ignored.

The *Pglob* parameter has meaning even if the **glob** subroutine fails. Therefore, the **glob** subroutine can report partial results in the event of an error. However, if the number of matched path names is 0, the pointer in the *Pglob* parameter is unspecified even if the **glob** subroutine did not return an error.

## Examples

The **GLOB\_NOCHECK** flag can be used with an application to expand any path name using wildcard characters. However, the **GLOB\_NOCHECK** flag treats the pattern as just a string by default. The **sh** command can use this facility for option parameters, for example.

The **GLOB\_DOOFFS** flag can be used by applications that build an argument list for use with the **execv**, **execve**, or **execvp** subroutine. For example, an application needs to do the equivalent of `ls -l *.c`, but for some reason cannot. The application could still obtain approximately the same result using the sequence:

```
globbuf.gl_offs = 2;
glob ("*.c", GLOB_DOOFFS, NULL, &globbuf);
globbuf.gl_pathv[0] = "ls";
globbuf.gl_pathv[1] = "-l";
execvp ("ls", &globbuf.gl_pathv[0]);
```

Using the same example, `ls -l *.c *.h` could be approximated using the **GLOB\_APPEND** flag as follows:

```
globbuf.gl_offs = 2;
glob ("*.c", GLOB_DOOFFS, NULL, &globbuf);
glob ("*.h", GLOB_DOOFFS|GLOB_APPEND, NULL, &globbuf);
```

The new path names generated by a subsequent call with the **GLOB\_APPEND** flag set are not sorted together with the previous path names. This is the same way the shell handles path name expansion when multiple expansions are done on a command line.

## Related Information

The **exec**: **execl**, **execv**, **execle**, **execve**, **execlp**, **execvp**, or **exec** (“exec: execl, execle, execlp, execv, execve, execvp, or exec Subroutine” on page 236) subroutine, **fnmatch** (“fnmatch Subroutine” on page 286) subroutine, **opendir**, **readdir**, **telldir**, **seekdir**, **rewinddir**, or **closedir** (“opendir, readdir, telldir,

seekdir, rewinddir, closedir, opendir64, readdir64, telldir64, seekdir64, rewinddir64, or closedir64 Subroutine” on page 943) subroutine, **statx**, **stat**, **lstat**, **fstatx**, **fstat**, **fullstat**, or **fullstat** subroutine.

The **ls** command.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## globfree Subroutine

### Purpose

Frees all memory associated with the *pglob* parameter.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <glob.h>
```

```
void globfree ( pglob )  
glob_t *pglob;
```

### Description

The **globfree** subroutine frees any memory associated with the *pglob* parameter due to a previous call to the **glob** subroutine.

### Parameters

*pglob*                      Structure containing the results of a previous call to the **glob** subroutine.

### Related Information

The **glob** (“glob Subroutine” on page 482) subroutine.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## grantpt Subroutine

### Purpose

Changes the mode and ownership of a pseudo-terminal device.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
int grantpt ( FileDescriptor )  
int FileDescriptor;
```

## Description

The **grantpt** subroutine changes the mode and the ownership of the slave pseudo-terminal associated with the master pseudo-terminal device defined by the *FileDescriptor* parameter. The user ID of the slave pseudo-terminal is set to the real UID of the calling process. The group ID of the slave pseudo-terminal is set to an unspecified group ID. The permission mode of the slave pseudo-terminal is set to readable and writeable by the owner, and writeable by the group.

## Parameters

*FileDescriptor* Specifies the file descriptor of the master pseudo-terminal device.

## Return Value

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **grantpt** function may fail if:

<b>EBADF</b>	The <i>fildev</i> argument is not a valid open file descriptor.
<b>EINVAL</b>	The <i>fildev</i> argument is not associated with a master pseudo-terminal device.
<b>EACCES</b>	The corresponding slave pseudo-terminal device could not be accessed.

## Related Information

The **unlockpt** subroutine.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## HBA\_CloseAdapter Subroutine

### Purpose

Closes the adapter opened by the **HBA\_OpenAdapter** subroutine.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
void HBA_CloseAdapter (handle)  
HBA_HANDLE handle;
```

### Description

The **HBA\_CloseAdapter** subroutine closes the file associated with the file handle that was the result of a call to the **HBA\_OpenAdapter** subroutine. The **HBA\_CloseAdapter** subroutine calls the **close** subroutine, and applies it to the file handle. After performing the operation, the handle is set to NULL.

### Parameters

*handle* Specifies the open file descriptor obtained from a successful call to the **open** subroutine.

## Related Information

The “HBA\_OpenAdapter Subroutine” on page 503.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_FreeLibrary Subroutine

### Purpose

Frees all the resources allocated to build the Common HBA API Library.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_FreeLibrary ()
```

### Description

The **HBA\_FreeLibrary** subroutine frees all resources allocated to build the Common HBA API library. This subroutine must be called after calling any other routine from the Common HBA API library.

### Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.

## Related Information

The “HBA\_LoadLibrary Subroutine” on page 502.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_GetAdapterAttributes, HBA\_GetPortAttributes, HBA\_GetDiscoveredPortAttributes, HBA\_GetPortAttributesByWWN Subroutine

### Purpose

Gets the attributes of the end device’s adapter, port, or remote port.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

## Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_GetAdapterAttributes (handle, hbaattributes)
HBA_STATUS HBA_GetAdapterPortAttributes (handle, portindex, portattributes)
HBA_STATUS HBA_GetDiscoveredPortAttributes (handle, portindex, discoveredportindex, portattributes)
HBA_STATUS HBA_GetPortAttributesByWWN (handle, PortWWN, portattributes)
```

```
HBA_HANDLE handle;
HBA_ADAPTERATTRIBUTES *hbaattributes;
HBA_UINT32 portindex;
HBA_PORTATTRIBUTES *portattributes;
HBA_UINT32 discoveredportindex;
HBA_WWN PortWWN;
```

## Description

The **HBA\_GetAdapterAttributes** subroutine queries the ODM and makes system calls to gather information pertaining to the adapter. This information is returned to the **HBA\_ADAPTERATTRIBUTES** structure. This structure is defined in the **/usr/include/sys/hbaapi.h** file.

The **HBA\_GetAdapterAttributes**, **HBA\_GetAdapterPortAttributes**, **HBA\_GetDiscoveredPortAttributes**, and **HBA\_GetPortAttributesByWWN** subroutines return the attributes of the adapter, port or remote port.

These attributes include:

- Manufacturer
- SerialNumber
- Model
- ModelDescription
- NodeWWN
- NodeSymbolicName
- HardwareVersion
- DriverVersion
- OptionROMVersion
- FirmwareVersion
- VendorSpecificID
- NumberOfPorts
- Drivename

The **HBA\_GetAdapterPortAttributes**, **HBA\_GetDiscoveredPortAttributes**, and **HBA\_GetPortAttributesByWWN** subroutines also query the ODM and make system calls to gather information. The gathered information pertains to the port attached to the adapter or discovered on the network. The attributes are stored in the **HBA\_PORTATTRIBUTES** structure. This structure is defined in the **/usr/include/sys/hbaapi.h** file.

These attributes include:

- NodeWWN
- PortWWN
- PortFcid
- PortType
- PortState
- PortSupportedClassofService
- PortSupportedFc4Types



- PortActiveFc4Types
- OSDeviceName
- PortSpeed
- NumberOfDiscoveredPorts
- PortSymbolicName
- PortSupportedSpeed
- PortMaxFrameSize
- FabricName

The **HBA\_GetAdapterPortAttributes** subroutine returns the attributes of the attached port.

The **HBA\_GetDiscoveredPortAttributes**, and **HBA\_GetPortAttributesByWWN** subroutines return the same information. However, these subroutines differ in the way they are called, and in the way they acquire the information.

## Parameters

<i>handle</i>	Specifies the open file descriptor obtained from a successful call to the <b>open</b> subroutine.
<i>hbaattributes</i>	Points to an <b>HBA_AdapterAttributes</b> structure, which is used to store information pertaining to the Host Bus Adapter.
<i>portindex</i>	Specifies the index number of the port where the information was obtained.
<i>portattributes</i>	Points to an <b>HBA_PortAttributes</b> structure used to store information pertaining to the port attached to the Host Bus Adapter.
<i>discoveredportindex</i>	Specifies the index of the attached port discovered over the network.
<i>PortWWN</i>	Specifies the world wide name or port name of the target device.

## Return Values

Upon successful completion, the attributes and a value of HBA\_STATUS\_OK, or 0 are returned.

If no information for a particular attribute is available, a null value is returned for that attribute. HBA\_STATUS\_ERROR or 1 is returned if certain ODM queries or system calls fail while trying to retrieve the attributes.

## Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.
<b>HBA_STATUS_ERROR_ARG</b>	A value of 4 if there was a bad argument.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	A value of 5 if the world wide name was not recognized.

## Related Information

“HBA\_GetAdapterName Subroutine” on page 490, and “HBA\_GetNumberOfAdapters Subroutine” on page 497.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_GetAdapterName Subroutine

### Purpose

Gets the name of a Common Host Bus Adapter.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_GetAdapterName (adapterindex, adaptername)
HBA_UINT32 adapterindex;
char *adaptername;
```

### Description

The **HBA\_GetAdapterName** subroutine gets the name of a Common Host Bus Adapter. The *adapterindex* parameter is an index into an internal table containing all FCP adapters on the machine. The *adapterindex* parameter is used to search the table and obtain the adapter name. The name of the adapter is returned in the form of `mgfdomain-model-adapterindex`. The name of the adapter is used as an argument for the **HBA\_OpenAdapter** subroutine. From the **HBA\_OpenAdapter** subroutine, the file descriptor will be obtained where additional Common HBA API routines can then be called using the file descriptor as the argument.

### Parameters

<i>adapterindex</i>	Specifies the index of the adapter held in the adapter table for which the name of the adapter is to be returned.
<i>adaptername</i>	Points to a character string that will be used to hold the name of the adapter.

### Return Values

Upon successful completion, the **HBA\_GetAdapterName** subroutine returns the name of the adapter and a 0, or a status code of `HBA_STATUS_OK`. If unsuccessful, a null value will be returned for *adaptername* and an value of 1, or a status code of `HBA_STATUS_ERROR`.

### Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_NOT_SUPPORTED</b>	A value of 2 if the function is not supported.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.
<b>HBA_STATUS_ERROR_ARG</b>	A value of 4 if there was a bad argument.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	A value of 5 if the world wide name was not recognized.
<b>HBA_STATUS_ERROR_ILLEGAL_INDEX</b>	A value of 6 if an index was not recognized.
<b>HBA_STATUS_ERROR_MORE_DATA</b>	A value of 7 if a larger buffer is required.
<b>HBA_STATUS_ERROR_STALE_DATA</b>	A value of 8 if information has changed since the last call to the <b>HBA_RefreshInformation</b> subroutine.

**HBA\_STATUS\_SCSI\_CHECK\_CONDITION**

A value of 9 if a SCSI Check Condition was reported.

**HBA\_STATUS\_ERROR\_BUSY**

A value of 10 if the adapter was busy or reserved. Try again later.

**HBA\_STATUS\_ERROR\_TRY\_AGAIN**

A value of 11 if the request timed out. Try again later.

**HBA\_STATUS\_ERROR\_UNAVAILABLE**

A value of 12 if the referenced HBA has been removed or deactivated.

## Related Information

The “HBA\_GetNumberOfAdapters Subroutine” on page 497.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_GetEventBuffer Subroutine

### Purpose

Removes and returns the next events from the HBA’s event queue.

### Syntax

```
HBA_STATUS HBA_GetEventBuffer(  
    HBA_HANDLE handle,  
    HBA_EVENTINFO *pEventBuffer,  
    HBA_UINT32 *pEventCount,  
);
```

### Description

The **HBA\_GetEventBuffer** function removes and returns the next events from the HBA’s event queue. The number of events returned is the lesser of the value of the *EventCount* parameter at the time of the call and the number of entries available in the event queue.

### Parameters

<i>handle</i>	A handle to an open HBA.
<i>pEventBuffer</i>	Pointer to a buffer to receive events.
<i>pEventCount</i>	Pointer to the number of event records that fit in the space allocated for the buffer to receive events. It is set to the size (in event records) of the buffer for receiving events on call, and is returned as the number of events actually delivered.

### Return Values

The value of the **HBA\_GetEventBuffer** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that no errors were encountered and *pEventCount* indicates the number of event records returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pEventBuffer</i>	Remains unchanged. The buffer to which it points contains event records representing previously undelivered events.
---------------------	---

*pEventCount*

Remains unchanged. The value of the integer to which it points contains the number of event records that actually were delivered.

## Error Codes

**HBA\_STATUS\_ERROR**

Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetFC4Statistics Subroutine,” “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_GetFC4Statistics Subroutine

### Purpose

Returns traffic statistics for a specific FC-4 protocol through a specific local HBA and local end port.

### Syntax

```
HBA_STATUS HBA_GetFC4Statistics(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_UINT8 FC4type,  
    HBA_FC4STATISTICS *statistics  
);
```

### Description

The **HBA\_GetFC4Statistics** function returns traffic statistics for a specific FC-4 protocol through a specific local HBA and local end port.

**Note:** Basic Link Service, Extended Link Service, and CT each have specific Data Structure **TYPE** values, so their traffic can be requested.

### Parameters

<i>handle</i>	A handle to an open HBA containing the end port for which FC-4 statistics can return.
<i>hbaPortWWN</i>	The Port Name of the local HBA end port for which FC-4 statistics can return.
<i>FC4type</i>	The Data Structure <b>TYPE</b> assigned by FC-FS to the FC-4 protocol for which FC-4 statistics are requested.
<i>statistics</i>	A pointer to an FC-4 Statistics structure in which the statistics for the specified FC-4 protocol can be returned.

### Return Values

The value of the **HBA\_GetFC4Statistics** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the statistics for the specified FC-4 and end port have been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return value for the following parameter is as follows:

*statistics*                Remains unchanged. The structure to which it points contains the statistics for the specified FC-4 protocol.

## Error Codes

<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	Indicates that the HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_UNSUPPORTED_FC4</b>	Indicates that the specified HBA end port does not support the specified FC-4 protocol.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_GetFcpPersistentBinding Subroutine

### Purpose

Gets persistent binding information of SCSI LUNs.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_GetFcpPersistentBinding (handle, binding)  
HBA_HANDLE handle;  
PHBA_FCPBinding binding;
```

### Description

For the specified HBA\_HANDLE, the **HBA\_GetFcpPersistentBinding** subroutine returns the full binding information of local SCSI LUNs to FCP LUNs for each child of the specified HBA\_HANDLE. Applications must allocate memory for the **HBA\_FCPBINDING** structure, and also pass to the subroutine the number of entries allocated. If the subroutine determines that the structure is not large enough to represent the full binding information, it will set the *NumberOfEntries* variable to the correct value and return an error.

### Parameters

*handle*                An HBA\_HANDLE to an open adapter.

*binding* A pointer to a structure containing the binding information of the handle's children. The **HBA\_FCPBINDING** structure has the following form:

```
struct HBA_FCPBinding {
    HBA_UINT32 NumberOfEntries;
    HBA_FCPBINDINGENTRY entry[1]; /* Variable length array */
};
```

The size of the structure is determined by the calling application, and is passed in by the *NumberOfEntries* variable.

## Return Values

Upon successful completion, **HBA\_STATUS\_OK** is returned, and the *binding* parameter points to the full binding structure. If the application has not allocated enough space for the full binding, **HBA\_STATUS\_ERROR\_MORE\_DATA** is returned and the *NumberOfEntries* field in the binding structure is set to the correct value.

## Error Codes

If there is insufficient space allocated for the full binding, **HBA\_STATUS\_ERROR\_MORE\_DATA** is returned.

## Related Information

The “HBA\_GetFcpTargetMapping Subroutine” on page 496.

---

## HBA\_GetFCPStatistics Subroutine

### Purpose

Returns traffic statistics for a specific OS SCSI logical unit provided by the FCP protocol on a specific local HBA.

### Syntax

```
HBA_STATUS HBA_GetFCPStatistics(
    HBA_HANDLE handle,
    const HBA_SCSIID *lunit,
    HBA_FC4STATISTICS *statistics
);
```

### Description

The **HBA\_GetFCPStatistics** function returns traffic statistics for a specific OS SCSI logical unit provided by the FCP protocol on a specific local HBA.

### Parameters

<i>handle</i>	A handle to an open HBA containing the end port for which FCP-2 statistics can be returned.
<i>lunit</i>	Pointer to a structure specifying the OS SCSI logical unit for which FCP-2 statistics are requested.
<i>statistics</i>	Pointer to a FC-4 Statistics structure in which the FCP-2 statistics for the specified logical unit can be returned.

### Return Values

The value of the **HBA\_GetFCPStatistics** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that FCP-2 statistics

have been returned for the specified HBA. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return value for the following parameter is as follows:

*statistics*                Remains unchanged. The structure to which it points contains the FCP-2 statistics for the specified HBA and logical unit.

## Error Codes

<b>HBA_STATUS_ERROR_INVALID_LUN</b>	The HBA referenced by <i>handle</i> does not support the logical unit referenced by <i>lunit</i> .
<b>HBA_STATUS_ERROR_UNSUPPORTED_FC4</b>	The specified HBA end port does not support FCP-2.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFcpTargetMappingV2 Subroutine,” “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_GetFcpTargetMappingV2 Subroutine

### Purpose

Returns the mapping between OS targets or logical units and FCP targets or logical units offered by the specified HBA end port at the time the function call is processed.

### Syntax

```
HBA_STATUS HBA_GetFcpTargetMappingV2(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_FCPTARGETMAPPINGV2 *pMapping  
);
```

### Description

The **HBA\_GetFcpTargetMappingV2** function returns the mapping between OS identification of SCSI targets or logical units and FCP identification of targets or logical units offered by the specified HBA end port at the time the function call is processed. Space in the *pMapping* structure permitting, one mapping entry is returned for each FCP logical unit represented in the OS and one mapping entry is returned for each FCP target that is represented in the OS but for which no logical units are represented in the OS. No target mapping entries are returned to represent FCP objects that are not represented in the OS (that is, objects that are unmapped).

The mappings returned include a Logical Unit Unique Device Identifier (LUID) for each logical unit that provides one. For logical units that provide more than one LUID, the LUID returned is the type 3 (FC **Name\_Identifier**) LUID with the smallest identifier value if any LUID of type 3 is provided; otherwise, the type 2 (IEEE EUI-64) LUID with the smallest identifier value if any LUID of type 2 is provided; otherwise, the type 1 (T10 vendor identification) LUID with the smallest identifier value if any LUID of type 1 is provided; otherwise, the type 0 (vendor specific) LUID with the smallest identifier value. If the logical unit provides no LUID, the value of the first four bytes of the LUID field are 0.

## Parameters

<i>handle</i>	A handle to an open HBA containing the end port for which target mappings are requested.
<i>hbaPortWWN</i>	Port Name of the local HBA end port for which target mappings are requested.
<i>pMapping</i>	Pointer to an <b>HBA_FCPTARGETMAPPINGV2</b> structure. The size of this structure shall be limited by the <i>NumberOfEntries</i> value within the structure.

## Return Values

The value of the **HBA\_GetFcpTargetMappingV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that all mapping entries have been returned for the specified end port. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return value for the following parameter is as follows:

<i>pMapping</i>	Remains unchanged. The structure to which it points contains mapping information from OS identifications of SCSI logical units to FCP identifications of logical units for the specified local HBA end port. The number of entries in the structure is the minimum of the number of entries specified at function call or the full mapping. The value of the <i>NumberOfEntries</i> field of the returned structure is the total number of mappings the end port has established. This is true even when the function returns an error stating that the buffer is too small to return all of the established mappings. An upper-level application can either allocate a sufficiently large buffer and check this value after a read, or do a read of the <i>NumberOfEntries</i> value separately and allocate a new buffer given the size to accommodate the entire mapping structure.
-----------------	--

## Error Codes

<b>HBA_STATUS_ERROR_MORE_DATA</b>	More space in the buffer is required to contain mapping information.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_NOT_SUPPORTED</b>	The HBA referenced by <i>handle</i> does not support target mapping.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_GetFcpTargetMapping Subroutine

### Purpose

Gets mapping of OS identification to FCP identification for each child of the specified HBA\_HANDLE.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)



## Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_GetFcpTargetMapping (handle, mapping)
```

```
HBA_HANDLE handle;
```

```
PHBA_FCPTARGETMAPPING mapping;
```

## Description

For the specified HBA\_HANDLE, the **HBA\_GetFcpTargetMapping** subroutine maps OS identification of all its SCSI logical units to their FCP identification. Applications must allocate memory for the **HBA\_FCPTargetMapping** structure, and also pass to the subroutine the number of entries allocated. If the subroutine determines that the structure is not large enough to represent the entire mapping, it will set the *NumberOfEntries* variable to the correct value and return an error.

## Parameters

*handle* An HBA\_HANDLE to an open adapter.

*mapping* A pointer to a structure containing a mapping of the handle's children. The **HBA\_FCPTARGETMAPPING** structure has the following form:

```
struct HBA_FCPTargetMapping (  
  HBA_UINT32 NumberOfEntries;  
  HBA_FCPCSIENTRY entry[1] /* Variable length array containing mappings */  
);
```

The size of the structure is determined by the calling application, and is passed in by the *NumberOfEntries* variable.

## Return Values

If successful, HBA\_STATUS\_OK is returned and the mapping parameter points to the full mapping structure. If the application has not allocated enough space for the full mapping, HBA\_STATUS\_ERROR\_MORE\_DATA is returned, and the *NumberOfEntries* field in the mapping structure is set to the correct value.

## Error Codes

If there is insufficient space allocated for the full mapping, HBA\_STATUS\_ERROR\_MORE\_DATA is returned.

## Related Information

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_GetNumberOfAdapters Subroutine

### Purpose

Returns the number of adapters discovered on the system.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

## Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_UINT32 HBA_GetNumberOfAdapters ()
```

## Description

The **HBA\_GetNumberOfAdapters** subroutine returns the number of HBAs supported by the library. The value returned is the current number of HBAs and reflects dynamic change of the HBA inventory without requiring a restart of the system, driver, or library.

## Return Values

The **HBA\_GetNumberOfAdapters** subroutine returns an integer representing the number of adapters on the machine.

## Related Information

The “HBA\_GetAdapterName Subroutine” on page 490.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_GetPersistentBindingV2 Subroutine

### Purpose

Returns persistent bindings between an FCP target and a SCSI ID for a specified HBA end port.

### Syntax

```
HBA_STATUS HBA_GetPersistentBindingV2(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_FCPTARGETMAPPINGV2 *binding  
);
```

### Description

The **HBA\_GetFcpPersistentBindingV2** function returns persistent bindings between an FCP target and a SCSI ID for a specified HBA end port. The binding information can include bindings to Logical Unit Unique Device Identifiers (LUIDs).

### Parameters

<i>handle</i>	A handle to an open HBA containing the end port for which persistent binding can be returned.
<i>hbaPortWWN</i>	The Port Name of the local HBA end port for which persistent binding can be returned.
<i>binding</i>	Pointer to an <b>HBA_FCPBINDING2</b> structure. The <i>NumberOfEntries</i> field in the structure limits the number of entries that are returned.

### Return Values

The value of the **HBA\_GetPersistentBindingV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that all binding entries have been returned for the specified end port. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return value for the following parameter is as follows:

*binding* Remains unchanged. The structure to which it points contains binding information from OS identifications of SCSI logical units to FCP and LUID identifications of logical units for the specified HBA end port. The number of entries in the structure is the minimum of the number of entries specified at function call or the full set of bindings. The *NumberOfEntries* field contains the total number of bindings established by the end port. An application can either call **HBA\_GetPersistentBindingV2** with *NumberOfEntries* set to 0 to retrieve the number of entries available, or allocate a sufficiently large buffer to retrieve entries at first call. The **Status** field of each **HBA\_FCPBINDINGENTRY2** substructure is 0.

## Error Codes

<b>HBA_STATUS_ERROR_MORE_DATA</b>	More space in the buffer is required to contain binding information.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_NOT_SUPPORTED</b>	The HBA referenced by <i>handle</i> does not support persistent binding.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_GetPortStatistics Subroutine

### Purpose

Gets the statistics for a Host Bus Adapter (HBA).

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_GetPortStatistics (handle, portindex, portstatistics)
HBA_HANDLE handle;
HBA_UINT32 portindex;
HBA_PORTSTATISTICS *portstatistics;
```

### Description

The **HBA\_GetPortStatistics** subroutine retrieves the statistics for the specified adapter. Only single-port adapters are supported, and the *portindex* parameter is disregarded. The exact meaning of events being counted for each statistic is vendor specific. The **HBA\_PORTSTATISTICS** structure includes the following fields:

- *SecondsSinceLastReset*

- *TxFrames*
- *TxWords*
- *RxFrames*
- *RxWords*
- *LIPCount*
- *NOSCount*
- *ErrorFrames*
- *DumpedFrames*
- *LinkFailureCount*
- *LossOfSyncCount*
- *LossOfSignalCount*
- *PrimitiveSeqProtocolErrCount*
- *InvalidTxWordCount*
- *InvalidCRCCount*

## Parameters

<i>handle</i>	HBA_HANDLE to an open adapter.
<i>portindex</i>	Not used.
<i>portstatistics</i>	Pointer to an HBA_PORTSTATISTICS structure.

## Return Values

Upon successful completion, HBA\_STATUS\_OK is returned. If the subroutine is unable to retrieve the statistics for an HBA, it returns HBA\_STATUS\_ERROR.

---

## HBA\_GetRNIDMgmtInfo Subroutine

### Purpose

Sends a **SCSI GET RNID** command to a remote port of the end device.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_GetRNIDMgmtInfo (handle, pInfo)
HBA_HANDLE handle;
HBA_MGMTINFO *pInfo;
```

### Description

The **HBA\_SetRNIDMgmtInfo** subroutine sends a **SCSI GET RNID** (Request Node Identification Data) command through a call to **ioctl** with the **SCIOCHBA** operation as its argument. The *arg* parameter for the **SCIOCHBA** operation is the address of a **scsi\_chba** structure. This structure is defined in the **/usr/include/sys/scsi\_buf.h** file. The *scsi\_chba* parameter block allows the caller to select the **GET RNID** command to be sent to the adapter. The **plinfo** structure stores the RNID data returned from **SCIOCHBA**. The **plinfo** structure is defined in the **/usr/include/sys/hbaapi.h** file. The structure includes:

- *wwn*
- *unittyp*

- PortId
- NumberOfAttachedNodes
- IPVersion
- UDPPort
- IPAddress
- reserved
- TopologyDiscoveryFlags

If successful, the GET RNID data in *pInfo* is returned from the adapter.

## Parameters

*handle* Specifies the open file descriptor obtained from a successful call to the **open** subroutine.  
*pInfo* Specifies the structure containing the information to get or set from the **RNID** command

## Return Values

Upon successful completion, the **HBA\_GetRNIDMgmtInfo** subroutine returns a pointer to a structure containing the data from the **GET RNID** command and a value of HBA\_STATUS\_OK, or a value of 0. If unsuccessful, a null value is returned along with a value of HBA\_STATUS\_ERROR, or a value of 1.

Upon successful completion, the **HBA\_SetRNIDMgmtInfo** subroutine returns a value of HBA\_STATUS\_OK, or a value of 0. If unsuccessful, an HBA\_STATUS\_ERROR value, or a value of 1 is returned.

## Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.

## Related Information

“HBA\_SendScsiInquiry Subroutine” on page 521, “HBA\_SendReadCapacity Subroutine” on page 512, “HBA\_SendCTPassThru Subroutine” on page 510, “HBA\_SendReportLUNs Subroutine” on page 513, “HBA\_SendRNID Subroutine” on page 516, and “HBA\_SetRNIDMgmtInfo Subroutine” on page 523.

SCSI Adapter Device Driver in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 2*.

Special Files in *AIX 5L Version 5.3 Files Reference*.

SCSI Subsystem Overview, A Typical Initiator-Mode SCSI Driver Transaction Sequence, Required SCSI Adapter Device Driver ioctl Commands, Understanding the Execution of Initiator I/O Requests, SCSI Error Recovery, and Understanding the *sc\_buf* Structure in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

---

## HBA\_GetVersion Subroutine

### Purpose

Returns the version number of the Common HBA API.

## Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

## Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_UINT32 HBA_GetVersion ()
```

## Description

The **HBA\_GetVersion** subroutine returns the version number representing the release of the Common HBA API.

## Return Values

Upon successful completion, the **HBA\_GetVersion** subroutine returns an integer value designating the version number of the Common HBA API.

## Related Information

“HBA\_LoadLibrary Subroutine” and “HBA\_FreeLibrary Subroutine” on page 487.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_LoadLibrary Subroutine

### Purpose

Loads a vendor specific library from the Common HBA API.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_LoadLibrary ()
```

### Description

The **HBA\_LoadLibrary** subroutine loads a vendor specific library from the Common HBA API. This library must be called first before calling any other routine from the Common HBA API.

### Return Values

The **HBA\_LoadLibrary** subroutine returns a value of 0, or **HBA\_STATUS\_OK**.

### Related Information

The “HBA\_FreeLibrary Subroutine” on page 487.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_OpenAdapter Subroutine

### Purpose

Opens the specified adapter for reading.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_HANDLE HBA_OpenAdapter (adaptername)  
char *adaptername;
```

### Description

The **HBA\_OpenAdapter** subroutine opens the adapter for reading for the purpose of getting it ready for additional calls from other subroutines in the Common HBA API.

The **HBA\_OpenAdapter** subroutine allows an application to open a specified HBA device, giving the application access to the device through the HBA\_HANDLE return value. The library ensures that all access to this HBA\_HANDLE between **HBA\_OpenAdapter** and **HBA\_CloseAdapter** calls is to the same device.

### Parameters

*adaptername* Specifies a string that contains the description of the adapter as returned by the **HBA\_GetAdapterName** subroutine.

### Return Values

If successful, the **HBA\_OpenAdapter** subroutine returns an HBA\_HANDLE with a value greater than 0. If unsuccessful, the subroutine returns a 0.

### Related Information

“HBA\_CloseAdapter Subroutine” on page 486, and “HBA\_GetAdapterName Subroutine” on page 490.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_OpenAdapterByWWN Subroutine

### Purpose

Attempts to open a handle to the HBA that contains a **Node\_Name** or **N\_Port\_Name** matching the *wwn* argument.

### Syntax

```
HBA_STATUS HBA_OpenAdapterByWWN(  
    HBA_HANDLE *pHandle,  
    HBA_WWN wwn  
);
```

## Description

The **HBA\_OpenAdapterByWWN** function attempts to open a handle to the HBA that contains a **Node\_Name** or **N\_Port\_Name** matching the *wwn* argument. The specified **Name\_Identifier** matches the **Node\_Name** or **N\_Port\_Name** of the HBA. Discovered end ports (remote end ports) are *not* checked for a match.

## Parameters

*pHandle* Pointer to a handle. The value at entry is irrelevant.  
*wwn* **Name\_Identifier** to match the **Node\_Name** or **N\_Port\_Name** of the HBA to open.

## Return Values

The value of the **HBA\_OpenAdapterByWWN** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the handle contains a valid HBA handle.

The return values for the following parameter is as follows:

*pHandle* Remains unchanged. If the open succeeds, the value to which it points is a handle to the requested HBA. On failure, the value is undefined.

## Error Codes

<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	There is no HBA with a <b>Node_Name</b> or <b>N_Port_Name</b> that matches <i>wwn</i> .
<b>HBA_STATUS_ERROR_AMBIGUOUS_WWN</b>	Multiple HBAs have a matching <b>Name_Identifier</b> . This can occur if the <b>Node_Names</b> of multiple HBAs are identical.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any other problem with opening the HBA.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_RefreshInformation Subroutine

### Purpose

Refreshes stale information from the Host Bus Adapter.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)



## Syntax

```
#include <sys/hbaapi.h>
```

```
void HBA_RefreshInformation (handle)  
HBA_HANDLE handle;
```

## Description

The **HBA\_RefreshInformation** subroutine refreshes stale information from the Host Bus Adapter. This would reflect changes to information obtained from calls to the **HBA\_GetAdapterPortAttributes**, or **HBA\_GetDiscoveredPortAttributes** subroutine. Once the application calls the **HBA\_RefreshInformation** subroutine, it can proceed to the attributes's call to get the new data.

## Parameters

*handle* Specifies the open file descriptor obtained from a successful call to the **open** subroutine for which the refresh operation is to be performed.

## Related Information

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_ScsiInquiryV2 Subroutine

### Purpose

Sends a SCSI INQUIRY command to a remote end port.

### Syntax

```
HBA_STATUS HBA_ScsiInquiryV2 (  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN discoveredPortWWN,  
    HBA_UINT64 fcLUN,  
    HBA_UINT8 CDB_Byte1,  
    HBA_UINT8 CDB_Byte2,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
    HBA_UINT8 *pScsiStatus,  
    void *pSenseBuffer,  
    HBA_UINT32 *pSenseBufferSize  
);
```

### Description

The **HBA\_ScsiInquiryV2** function sends a **SCSI INQUIRY** command to a remote end port.

A SCSI command is never sent to an end port that does not have SCSI target functionality. If sending a SCSI command causes a SCSI overlapped command condition with a correctly operating target, the command does not get sent. Proper use of tagged commands is an acceptable means of avoiding a SCSI overlapped command condition with targets that support tagged commands.

### Parameters

*handle* Open HBA through which the **SCSI INQUIRY** command can be issued.  
*hbaPortWWN* The Port Name for a local HBA end port through which the **SCSI INQUIRY** command can be issued.

<i>discoveredPortWWN</i>	The Port Name for an end port to which the <b>SCSI INQUIRY</b> command can be sent.
<i>fcLUN</i>	The SCSI LUN to which the <b>SCSI INQUIRY</b> command can be sent.
<i>CDB_Byte1</i>	The second byte of the CDB for the <b>SCSI INQUIRY</b> command. This contains control flag bits. At the time this standard was written, the effects of the value of <i>CDB_Byte1</i> on a SCSI INQUIRY command were as follows: <ul style="list-style-type: none"> <li>• 0 <ul style="list-style-type: none"> <li>– Requests the standard SCSI INQUIRY data.</li> </ul> </li> <li>• 1 <ul style="list-style-type: none"> <li>– Requests the vital product data (EVPD) specified by <i>CDB_Byte2</i>.</li> </ul> </li> <li>• 2 <ul style="list-style-type: none"> <li>– Requests command support data (<b>CmdDt</b>) for the command specified in <i>CDB_Byte2</i>.</li> </ul> </li> <li>• Other values <ul style="list-style-type: none"> <li>– Can cause SCSI Check Condition.</li> </ul> </li> </ul>
<i>CDB_Byte2</i>	The third byte of the CDB for the <b>SCSI INQUIRY</b> command. If <i>CDB_Byte1</i> is 1, <i>CDB_Byte2</i> is the Vital Product Data page code to request. If <i>CDB_Byte1</i> is 2, <i>CDB_Byte2</i> is the Operation Code of the command support data requested. For other values of <i>CDB_Byte1</i> , the value of <i>CDB_Byte2</i> is unspecified, and values other than 0 can cause a SCSI Check Condition.
<i>pRspBuffer</i>	A pointer to a buffer to receive the <b>SCSI INQUIRY</b> command response.
<i>pRspBufferSize</i>	A pointer to the size in bytes of the buffer to receive the <b>SCSI INQUIRY</b> command response.
<i>pScsiStatus</i>	A pointer to a buffer to receive SCSI status.
<i>pSenseBuffer</i>	A pointer to a buffer to receive SCSI sense data.
<i>pSenseBufferSize</i>	A pointer to the size in bytes of the buffer to receive SCSI sense data.

## Return Values

The value of the **HBA\_ScsilquiryV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete payload of a reply to the **SCSI INQUIRY** command has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. If the function value is <b>HBA_STATUS_OK</b> , the buffer to which it points contains the response to the <b>SCSI INQUIRY</b> command.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points is the size in bytes of the response returned by the command. This cannot exceed the size passed as an argument at this pointer.
<i>pScsiStatus</i>	Remains unchanged. The value of the byte to which it points is the SCSI status. If the function value is <b>HBA_STATUS_OK</b> or <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> , the value of the SCSI status can be interpreted based on the SCSI spec. A SCSI status of <b>HBA_STATUS_OK</b> indicates that a SCSI response is in the response buffer. A SCSI status of <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> indicates that no value is stored in the response, and the sense buffer contains failure information if available.
<i>pSenseBuffer</i>	Remains unchanged. If the function value is <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> , the buffer to which it points contains the sense data for the command.
<i>pSenseBufferSize</i>	Remains unchanged. The value of the integer to which it points is the size in bytes of the sense information returned by the command. This cannot exceed the size passed as an argument at this pointer.

## Error Codes

<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by handle does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_NOT_A_TARGET</b>	The identified remote end port does not have SCSI Target functionality.
<b>HBA_STATUS_ERROR_TARGET_BUSY</b>	Unable to send the requested command without causing a SCSI overlapped command condition.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiReadCapacityV2 Subroutine,” “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_ScsiReadCapacityV2 Subroutine

### Purpose

Sends a SCSI READ CAPACITY command to a remote end port.

### Syntax

```
HBA_STATUS HBA_ScsiReadCapacityV2(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN discoveredPortWWN,  
    HBA_UINT64 fcLUN,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
    HBA_UINT8 *pScsiStatus,  
    void *pSenseBuffer,  
    HBA_UINT32 *pSenseBufferSize  
);
```

### Description

The **HBA\_ScsiReadCapacityV2** function sends a SCSI READ CAPACITY command to a remote end port.

A SCSI command is never sent to an end port that does not have SCSI target functionality. If sending a SCSI command causes a SCSI overlapped command condition with a correctly operating target, the command will not be sent. Proper use of tagged commands is an acceptable means of avoiding a SCSI overlapped command condition with targets that support tagged commands.

### Parameters

<i>handle</i>	A handle to an open HBA through which the SCSI READ CAPACITY command is issued.
<i>hbaPortWWN</i>	The Port Name for a local HBA end port through which the SCSI READ CAPACITY command is issued.
<i>discoveredPortWWN</i>	The Port Name for an end port to which the SCSI READ CAPACITY command is sent.
<i>fcLUN</i>	The SCSI LUN to which the SCSI READ CAPACITY command is sent.
<i>pRspBuffer</i>	Pointer to a buffer to receive the SCSI READ CAPACITY command response.

<i>pRspBufferSize</i>	Pointer to the size in bytes of the buffer to receive the SCSI READ CAPACITY command response.
<i>pScsiStatus</i>	Pointer to a buffer to receive SCSI status.
<i>pSenseBuffer</i>	Pointer to a buffer to receive SCSI sense data.
<i>pSenseBufferSize</i>	Pointer to the size in bytes of the buffer to receive SCSI sense data.

## Return Values

The value of the **HBA\_ScsiReadCapacityV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete payload of a reply to the SCSI READ CAPACITY command has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. If the function value is <b>HBA_STATUS_OK</b> , the buffer to which it points contains the response to the SCSI READ CAPACITY command.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points is the size in bytes of the response returned by the command. This cannot exceed the size passed as an argument at this pointer.
<i>pScsiStatus</i>	Remains unchanged. The value of the byte to which it points is the SCSI status. If the function value is <b>HBA_STATUS_OK</b> or <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> , the value of the SCSI status can be interpreted based on the SCSI spec. A SCSI status of <b>HBA_STATUS_OK</b> indicates that a SCSI response is in the response buffer. A SCSI status of <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> indicates that no value is stored in the response, and the sense buffer contains failure information if available.
<i>pSenseBuffer</i>	Remains unchanged. If the function value is <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> , the buffer to which it points contains the sense data for the command.
<i>pSenseBufferSize</i>	Remains unchanged. The value of the integer to which it points is the size in bytes of the sense information returned by the command. This cannot exceed the size passed as an argument at this pointer.

## Error Codes

<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_NOT_A_TARGET</b>	The identified remote end port does not have SCSI Target functionality.
<b>HBA_STATUS_ERROR_TARGET_BUSY</b>	Unable to send the requested command without causing a SCSI overlapped command condition.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_ScsiReportLunsV2 Subroutine

### Purpose

Sends a SCSI REPORT LUNS command to Logical Unit Number 0 of a remote end port.

### Syntax

```
HBA_STATUS HBA_ScsiReportLUNsv2(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN discoveredPortWWN,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
    HBA_UINT8 *pScsiStatus,  
    void *pSenseBuffer,  
    HBA_UINT32 *pSenseBufferSize  
);
```

### Description

The **HBA\_ScsiReportLunsV2** function shall send a SCSI REPORT LUNS command to Logical Unit Number 0 of a remote end port.

A SCSI command is never sent to an end port that does not have SCSI target functionality. If sending a SCSI command causes a SCSI overlapped command condition with a correctly operating target, the command will not be sent. Proper use of tagged commands is an acceptable means of avoiding a SCSI overlapped command condition with targets that support tagged commands.

### Parameters

<i>handle</i>	A handle to an open HBA through which the SCSI REPORT LUNS command is issued.
<i>hbaPortWWN</i>	The Port Name for a local HBA end port through which the SCSI REPORT LUNS command is issued.
<i>discoveredPortWWN</i>	The Port Name for an end port to which the SCSI REPORT LUNS command is sent.
<i>pRspBuffer</i>	Pointer to a buffer to receive the SCSI REPORT LUNS command response.
<i>pRspBufferSize</i>	Pointer to the size in bytes of the buffer to receive the SCSI REPORT LUNS command response.
<i>pScsiStatus</i>	Pointer to a buffer to receive SCSI status.
<i>pSenseBuffer</i>	Pointer to a buffer to receive SCSI sense data.
<i>pSenseBufferSize</i>	Pointer to the size in bytes of the buffer to receive SCSI sense data.

### Return Values

The value of the **HBA\_ScsiReportLunsV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete payload of a reply to the SCSI REPORT LUNS command has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. If the function value is <b>HBA_STATUS_OK</b> , the buffer to which it points contains the response to the SCSI REPORT LUNS command.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points is the size in bytes of the response returned by the command. This cannot exceed the size passed as an argument at this pointer.

<i>pScsiStatus</i>	Remains unchanged. The value of the byte to which it points is the SCSI status. If the function value is <b>HBA_STATUS_OK</b> or <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> , the value of the SCSI status can be interpreted based on the SCSI spec. A SCSI status of <b>HBA_STATUS_OK</b> indicates that a SCSI response is in the response buffer. A SCSI status of <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> indicates that no value is stored in the response, and the sense buffer contains failure information if available.
<i>pSenseBuffer</i>	Remains unchanged. If the function value is <b>HBA_STATUS_SCSI_CHECK_CONDITION</b> , the buffer to which it points contains the sense data for the command.
<i>pSenseBufferSize</i>	Remains unchanged. The value of the integer to which it points is the size in bytes of the sense information returned by the command. This cannot exceed the size passed as an argument at this pointer.

## Error Codes

<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_NOT_A_TARGET</b>	The identified remote end port does not have SCSI Target functionality.
<b>HBA_STATUS_ERROR_TARGET_BUSY</b>	Unable to send the requested command without causing a SCSI overlapped command condition.
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_SendCTPassThru Subroutine

### Purpose

Sends a CT<sup>®</sup> pass through frame.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_SendCTPassThru (handle, pReqBuffer, ReqBufferSize, pRspBuffer, RspBufferSize)
HBA_HANDLE handle;
void *pReqBuffer;
HBA_UINT32 ReqBufferSize;
void *pRspBuffer;
HBA_UINT32 RspBufferSize;
```

### Description

The **HBA\_SendCTPassThru** subroutine sends a CT pass through frame to a fabric connected to the specified handle. The CT frame is routed in the fabric according to the *GS\_TYPE* field in the CT frame.

## Parameters

<i>handle</i>	HBA_HANDLE to an open adapter.
<i>pReqBuffer</i>	Pointer to a buffer that contains the CT request.
<i>ReqBufferSize</i>	Size of the request buffer.
<i>pRspBuffer</i>	Pointer to a buffer that receives the response of the command.
<i>RspBufferSize</i>	Size of the response buffer.

## Return Values

If successful, HBA\_STATUS\_OK is returned, and the *pRspBuffer* parameter points to the CT response.

## Error Codes

If the adapter specified by the *handle* parameter is connected to an arbitrated loop, the **HBA\_SendCTPassThru** subroutine returns HBA\_STATUS\_ERROR\_NOT\_SUPPORTED. This subroutine is only valid when connected to a fabric.

## Related Information

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_SendCTPassThruV2 Subroutine

### Purpose

Sends a CT request payload.

### Syntax

```
HBA_STATUS HBA_SendCTPassThruV2(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    void *pReqBuffer,  
    HBA_UINT32 *ReqBufferSize,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
);
```

### Description

The **HBA\_SendCTPassThruV2** function sends a CT request payload. An HBA should decode this CT\_IU request by, routing the CT frame in a fabric according to the **GS\_TYPE** field within the CT frame.

### Parameters

<i>handle</i>	A handle to an open HBA through which the CT request is issued.
<i>hbaPortWWN</i>	The Port Name for a local HBA <b>Nx_Port</b> through which the CT request is issued.
<i>pReqBuffer</i>	Pointer to a buffer containing the full CT payload, including the CT header, to be sent with byte ordering.
<i>ReqBufferSize</i>	The size of the full CT payload, including the CT header, in bytes.
<i>pRSPBuffer</i>	Pointer to a buffer for the CT response.
<i>pRSPBufferSize</i>	Pointer to the size in bytes of the buffer for the CT response payload.

## Return Values

The value of the **SendCTPassThruV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete reply to the CT **Passthru** command has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. The buffer to which it points contains the CT response payload, including the CT header received in response to the frame sent, with byte ordering. If the size of the actual response exceeds the size of the response buffer, trailing data is truncated from the response so that the returned data equals the size of the buffer.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points is set to the size (in bytes) of the actual response data.

## Error Codes

<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an <b>Nx_Port</b> with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_SendReadCapacity Subroutine

### Purpose

Sends a **SCSI READ CAPACITY** command to a Fibre Channel port.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_SendReadCapacity (handle, portWWN, fclUN, pRspBuffer, RspBufferSize, pSenseBuffer, SenseBufferSize)
HBA_HANDLE handle;
HBA_WWN portWWN;
HBA_UINT64 fclUN;
void *pRspBuffer;
HBA_UINT32 RspBufferSize;
void *pSenseBuffer;
HBA_UINT32 SenseBufferSize;
```



## Description

The **HBA\_SendReadCapacity** subroutine sends a **SCSI READ CAPACITY** command to the Fibre Channel port connected to the *handle* parameter and specified by the *portWWN* and *fcLUN* parameters.

## Parameters

<i>handle</i>	HBA_HANDLE to an open adapter.
<i>portWWN</i>	Port world-wide name of an adapter.
<i>fcLUN</i>	Fibre Channel LUN to send the <b>SCSI READ CAPACITY</b> command to.
<i>pRspBuffer</i>	Pointer to a buffer that receives the response of the command.
<i>RspBufferSize</i>	Size of the response buffer.
<i>pSenseBuffer</i>	Pointer to a buffer that receives sense information.
<i>SenseBufferSize</i>	Size of the sense buffer.

## Return Values

If successful, **HBA\_STATUS\_OK** is returned and the *pRspBuffer* parameter points to the response to the **READ CAPACITY** command. If an error occurs, **HBA\_STATUS\_ERROR** is returned.

## Error Codes

If the *portWWN* value is not a valid world-wide name connected to the specified handle, **HBA\_STATUS\_ERROR\_ILLEGAL\_WWN** is returned. On any other types of failures, **HBA\_STATUS\_ERROR** is returned.

## Related Information

The “HBA\_SendScsilnquiry Subroutine” on page 521.

Special Files in *AIX 5L Version 5.3 Files Reference* describes specific qualities of the files that define devices.

---

## HBA\_SendReportLUNs Subroutine

### Purpose

Sends a **SCSI REPORT LUNs** command to a remote port of the end device.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_SendReportLUNs (handle, PortWWN, pRspBuffer, RspBufferSize, pSenseBuffer, SenseBufferSize)
HBA_HANDLE handle;
HBA_WWN PortWWN;
void *pRspBuffer;
HBA_UINT32 RspBufferSize;
void *pSenseBuffer;
HBA_UINT32 SenseBufferSize;
```

### Description

The **HBA\_SendReportLUNs** subroutine sends a **SCSI REPORT LUNs** command through a call to **ioctl** with the **SCIOLCMD** operation as its argument. The *arg* parameter for the **SCIOLCMD** operation is the address of a **scsi\_iocmd** structure. This structure is defined in the **/usr/include/sys/scsi\_buf.h** file. The

*scsi\_iocmd* parameter block allows the caller to select the SCSI and LUN IDs to be queried. The caller also specifies the SCSI command descriptor block area, command length (SCSI command block length), the time-out value for the command, and a *flags* field.

If successful, the report LUNs data is returned in *pRspBuffer*. The returned report LUNs data must be examined to see if the requested LUN exists.

## Parameters

<i>handle</i>	Specifies the open file descriptor obtained from a successful call to the <b>open</b> subroutine.
<i>PortWWN</i>	Specifies the world wide name or port name of the target device.
<i>pRspBuffer</i>	Points to a buffer containing the requested instruction for a send/read capacity request to an open adapter.
<i>RspBufferSize</i>	Specifies the size of the buffer to the <i>pRspBuffer</i> parameter.
<i>pSenseBuffer</i>	Points to a buffer containing the data returned from a send/read capacity request to an open adapter.
<i>SenseBufferSize</i>	Specifies the size of the buffer to the <i>pSenseBuffer</i> parameter.

## Return Values

Upon successful completion, the **HBA\_SendReportLUNs** subroutine returns a buffer in bytes containing the SCSI report of LUNs, a buffer containing the SCSI sense data, and a value of HBA\_STATUS\_OK, or a value of 0.

If unsuccessful, an empty buffer for the SCSI report of LUNs, a response buffer containing the failure, and a value of HBA\_STATUS\_ERROR, or a value of 1 is returned.

## Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	A value of 5 if the world wide name was not recognized.
<b>HBA_STATUS_SCSI_CHECK_CONDITION</b>	A value of 9 if a SCSI Check Condition was reported.

## Related Information

“HBA\_SendScsiInquiry Subroutine” on page 521, “HBA\_SendReadCapacity Subroutine” on page 512, “HBA\_SendCTPassThru Subroutine” on page 510, “HBA\_SendRNID Subroutine” on page 516, “HBA\_SetRNIDMgmtInfo Subroutine” on page 523, and “HBA\_GetRNIDMgmtInfo Subroutine” on page 500.

SCSI Adapter Device Driver in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 2*.

Special Files in *AIX 5L Version 5.3 Files Reference*.

SCSI Subsystem Overview, A Typical Initiator-Mode SCSI Driver Transaction Sequence, Required SCSI Adapter Device Driver ioctl Commands, Understanding the Execution of Initiator I/O Requests, SCSI Error Recovery, and Understanding the *sc\_buf* Structure in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

---

## HBA\_SendRLS Subroutine

### Purpose

Issues a Read Link Error Status Block (RLS) Extended Link Service through the specified HBA end port.

### Syntax

```
HBA_STATUS HBA_SendRLS (  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN destWWN,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
);
```

### Description

The **HBA\_SendRLS** function issues a Read Link Error Status Block (RLS) Extended Link Service through the specified HBA end port to request a specified remote FC\_Port to return the Link Error Status Block associated with the destination Port Name.

### Parameters

<i>handle</i>	A handle to an open HBA through which the ELS is sent.
<i>hbaPortWWN</i>	Port Name of the local HBA end port through which the ELS is sent.
<i>destWWN</i>	Port Name of the remote FC_Port to which the ELS is sent.
<i>pRspBuffer</i>	Pointer to a buffer to receive the ELS response.
<i>pRSPBufferSize</i>	Pointer to the size in bytes of <i>pRspBuffer</i> . A size of 28 is sufficient for the largest response.

### Return Values

The value of the **HBA\_SendRLS** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete LS\_ACC to the RLS ELS has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. The buffer to which it points contains the payload data from the RLS Reply. Note that if the ELS was rejected, this is the LS_RJT payload. If the size of the reply payload exceeds the size specified in the <i>pRspBufferSize</i> parameter at entry to the function, the returned data is truncated to the size specified in the argument.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points contains the size in bytes of the complete ELS reply payload. This can exceed the size specified as an argument. This indicates that the data in <i>pRspBuffer</i> has been truncated.

### Error Codes

<b>HBA_STATUS_ERROR_ELS_REJECT</b>	The RNID ELS was rejected by the destination FC_Port.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_SendRNID Subroutine

### Purpose

Sends an RNID command through a call to **SCIOLPAYLD** to a remote port of the end device.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_SendRNID (handle, wwn, wwntype, pRspBuffer, RspBufferSize)
HBA_HANDLE handle;
HBA_WWN wwn;
HBA_WWNTYPE wwntype;
void *pRspBuffer;
HBA_UINT32 RspBufferSize;
```

### Description

The **HBA\_SendRNID** subroutine sends a **SCSI RNID** command with the Node Identification Data Format set to indicate the default Topology Discovery format. This is done through a call to **ioctl** with the **SCIOLPAYLD** operation as its argument. The *arg* parameter for the **SCIOLPAYLD** operation is the address of an **scsi\_trans\_payld** structure. This structure is defined in the **/usr/include/sys/scsi\_buf.h** file. The *scsi\_trans\_payld* parameter block allows the caller to select the SCSI and LUN IDs to be queried. In addition, the caller must specify the **fcph\_rnid\_payld\_t** structure to hold the command and the topology format for **SCIOLPAYLD**. The structure for the **fcph\_rnid\_payld\_t** structure is defined in the **/usr/include/sys/fcph.h** file.

If successful, the RNID data is returned in *pRspBuffer*. The returned RNID data must be examined to see if the requested information exists.

### Parameters

<i>handle</i>	Specifies the open file descriptor obtained from a successful call to the <b>open</b> subroutine.
<i>wwn</i>	Specifies the world wide name or port name of the target device.
<i>wwntype</i>	Specifies the type of the world wide name or port name of the target device.
<i>pRspBuffer</i>	Points to a buffer containing the requested instruction for a send/read capacity request to an open adapter.
<i>RspBufferSize</i>	Specifies the size of the buffer to the <i>pRspBuffer</i> parameter.

### Return Values

Upon successful completion, the **HBA\_SendRNID** subroutine returns a buffer in bytes containing the SCSI RNID data and a value of **HBA\_STATUS\_OK**, or a value of 0. If unsuccessful, an empty buffer for the SCSI RNID and a value of **HBA\_STATUS\_ERROR**, or a value of 1 is returned.

## Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_NOT_SUPPORTED</b>	A value of 2 if the function is not supported.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	A value of 5 if the world wide name was not recognized.

## Related Information

“HBA\_SendScsiInquiry Subroutine” on page 521, “HBA\_SendReadCapacity Subroutine” on page 512, “HBA\_SendCTPassThru Subroutine” on page 510, “HBA\_SendReportLUNs Subroutine” on page 513, “HBA\_SetRNIDMgmtInfo Subroutine” on page 523, and “HBA\_GetRNIDMgmtInfo Subroutine” on page 500.

SCSI Adapter Device Driver in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 2*.

Special Files in *AIX 5L Version 5.3 Files Reference*.

SCSI Subsystem Overview, A Typical Initiator-Mode SCSI Driver Transaction Sequence, Required SCSI Adapter Device Driver ioctl Commands, Understanding the Execution of Initiator I/O Requests, SCSI Error Recovery, and Understanding the `sc_buf` Structure in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

---

## HBA\_SendRNIDV2 Subroutine

### Purpose

Issues an RNID ELS to another FC\_Port requesting a specified Node Identification Data Format.

### Syntax

```
HBA_STATUS HBA_SendRNIDV2(  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN destWWN,  
    HBA_UINT32 destFCID,  
    HBA_UINT32 NodeIdDataFormat,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
);
```

### Description

The **HBA\_SendRNIDV2** function issues an RNID ELS to another FC\_Port requesting a specified Node Identification Data Format.

The *destFCID* parameter can be set to allow the RNID ELS to be sent to an FC\_Port that might not be registered with the name server. If *destFCID* is set to `x'00 00 00'`, the parameter is ignored. If *destFCID* is not 0, the HBA API library verifies that the *destWWN/destFCID* pair match in order to limit visibility that can violate scoping mechanisms (such as soft zoning):

- If the *destWWN/destFCID* pair matches an entry in the discovered ports table, the RNID is sent.
- If there is no entry in the discovered ports table for the *destWWN* or *destFCID*, the RNID is sent.
- If there is an entry in the discovered ports table for the *destWWN*, but the *destFCID* does not match, then the request is rejected.

- On completion of the **HBA\_SendRNIDV2**, if the Common Identification Data Length is nonzero in the RNID response, the API library compares the **N\_Port\_Name** in the Common Identification Data of the RNID response with *destWWN* and fails the operation without returning the response data if they do not match. If the Common Identification Data Length is 0 in the RNID response, this test is omitted.

## Parameters

<i>handle</i>	A handle to an open HBA through which the ELS is sent.
<i>hbaPortWWN</i>	Port Name of the local HBA end port through which the ELS is sent.
<i>destWWN</i>	Port Name of the remote FC_Port to which the ELS is sent.
<i>destFCID</i>	Address identifier of the destination to which the ELS is sent if <i>destFCID</i> is nonzero. <i>destFCID</i> is ignored if <i>destFCID</i> is 0.
<i>NodeIDDataFormat</i>	Valid value for Node Identification Data Format.
<i>pRSPBuffer</i>	Pointer to a buffer to receive the ELS response.
<i>pRSPBufferSize</i>	Pointer to the size in bytes of <i>pRspBuffer</i> .

## Return Values

The value of the **HBA\_SendRNIDV2** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete LS\_ACC to the RNID ELS has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. The buffer to which it points contains the payload data from the RNID Reply. Note that if the ELS was rejected, this is the LS_RJT payload. If the size of the reply payload exceeds the size specified in the <i>pRspBufferSize</i> parameter at entry to the function, the returned data is truncated to the size specified in the argument.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points contains the size in bytes of the complete ELS reply payload. This can exceed the size specified as an argument. This indicates that the data in <i>pRspBuffer</i> has been truncated.

## Error Codes

<b>HBA_STATUS_ERROR_ELS_REJECT</b>	The RNID ELS was rejected by the destination end port.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR_ILLEGAL_FCID</b>	The <i>destWWN/destFCID</i> pair conflicts with a discovered Port Name/address identifier pair known by the HBA referenced by <i>handle</i> .
<b>HBA_STATUS_ERROR_ILLEGAL_FCID</b>	The <b>N_Port_Name</b> in the RNID response does not match the <i>destWWN</i> .
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRPL Subroutine” on page 519, “HBA\_SendRPS Subroutine” on page 520

---

## HBA\_SendRPL Subroutine

### Purpose

Issues a Read Port List (RPL) Extended Link Service through the specified HBA to a specified end port or domain controller.

### Syntax

```
HBA_STATUS HBA_SendRPL (  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN agent_wwn,  
    HBA_UINT32 agent_domain,  
    HBA_UINT32 portIndex,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
);
```

### Description

The **HBA\_SendRPL** function issues a Read Port List (RPL) Extended Link Service through the specified HBA to a specified end port or domain controller.

### Parameters

<i>handle</i>	A handle to an open HBA through which the ELS is sent.
<i>hbaPortWWN</i>	Port Name of the local HBA end port through which the ELS is sent.
<i>agent_wwn</i>	Port Name of an FC_Port that is requested to provide its list of FC_Ports if <i>agent_wwn</i> is nonzero. If <i>agent_wwn</i> is 0, it is ignored.
<i>agent_domain</i>	Domain number and the domain controller for that domain shall be the entity that shall be requested to provide its list of FC_Ports if <i>agent_wwn</i> is 0. If <i>agent_wwn</i> is nonzero, <i>agent_domain</i> is ignored.
<i>portIndex</i>	Index of the first FC_Port requested in the response list. <b>Note:</b> If the recipient has proper compliance, the index of the first FC_Port in the complete list maintained by the recipient of the request is 0.
<i>pRSPBuffer</i>	Pointer to a buffer to receive the ELS response.
<i>pRSPBufferSize</i>	Pointer to the size in bytes of <i>pRspBuffer</i> . <b>Note:</b> If the responding entity has proper compliance, it truncates the list in the response to the number of FC_Ports that fit.

### Return Values

The value of the **HBA\_SendRPL** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete LS\_ACC to the RPL ELS has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. The buffer to which it points contains the payload data from the RPL Reply. If the ELS was rejected, this is the LS_RJT payload. If the size of the reply payload exceeds the size specified in the <i>pRspBufferSize</i> parameter at entry to the function, the returned data is truncated to the size specified in the argument.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points contains the size in bytes of the complete ELS reply payload. This can exceed the size specified as an argument. This indicates that the data in <i>pRspBuffer</i> has been truncated. <b>Note:</b> Truncation is not necessary if the responding entity is of proper compliance.

## Error Codes

**HBA\_STATUS\_ERROR\_ELS\_REJECT**  
**HBA\_STATUS\_ERROR\_ILLEGAL\_WWN**

The RPL ELS was rejected by the destination end port.  
The HBA referenced by *handle* does not contain an end port with Port Name *hbaPortWWN*.

**HBA\_STATUS\_ERROR**

Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPS Subroutine”

---

## HBA\_SendRPS Subroutine

### Purpose

Issues a Read Port Status Block (RPS) Extended Link Service through the specified HBA to a specified FC\_Port or domain controller.

### Syntax

```
HBA_STATUS HBA_SendRPS (  
    HBA_HANDLE handle,  
    HBA_WWN hbaPortWWN,  
    HBA_WWN agent_wwn,  
    HBA_UINT32 agent_domain,  
    HBA_WWN object_wwn,  
    HBA_UINT32 object_port_number,  
    void *pRspBuffer,  
    HBA_UINT32 *pRspBufferSize,  
);
```

### Description

The **HBA\_SendRPS** function issues a Read Port Status Block (RPS) Extended Link Service through the specified HBA to a specified FC\_Port or domain controller.

### Parameters

<i>handle</i>	A handle to an open HBA through which the ELS is sent.
<i>hbaPortWWN</i>	Port Name of the local HBA end port through which the ELS is sent.
<i>agent_wwn</i>	Port Name of an FC_Port that is requested to provide Port Status if <i>agent_wwn</i> is nonzero. <i>agent_wwn</i> is ignored if its value is 0.
<i>agent_domain</i>	Domain number for the domain controller that is requested to provide Port status if <i>agent_wwn</i> is 0. <i>agent_domain</i> is ignored if <i>agent_wwn</i> is nonzero.
<i>object_wwn</i>	Port Name of an FC_Port for which Port Status is returned if <i>object_wwn</i> is nonzero. <i>object_wwn</i> is ignored if its value is 0.
<i>object_port_number</i>	Relative port number of the FC_Port for which Port Status is returned if <i>object_wwn</i> is 0. The relative port number is defined in a vendor-specific manner within the entity to which the request is sent. <i>object_port_number</i> is ignored if <i>object_wwn</i> is nonzero.
<i>pRspBuffer</i>	Pointer to a buffer to receive the ELS response.
<i>pRSPBufferSize</i>	Pointer to the size in bytes of <i>pRspBuffer</i> . A size of 56 is sufficient for the largest response.



## Return Values

The value of the **HBA\_SendRPS** function is a valid status return value that indicates the reason for completion of the requested function. **HBA\_STATUS\_OK** is returned to indicate that the complete LS\_ACC to the RPS ELS has been returned. A valid status return value that most closely describes the result of the function should be returned to indicate a reason with no required value.

The return values for the following parameters are as follows:

<i>pRspBuffer</i>	Remains unchanged. The buffer to which it points contains the payload data from the RPS Reply. If the ELS was rejected, this is the LS_RJT payload. If the size of the reply payload exceeds the size specified in the <i>pRspBufferSize</i> parameter at entry to the function, the returned data is truncated to the size specified in the argument.
<i>pRspBufferSize</i>	Remains unchanged. The value of the integer to which it points contains the size in bytes of the complete ELS reply payload. This can exceed the size specified as an argument. This indicates that the data in <i>pRspBuffer</i> has been truncated.

## Error Codes

<b>HBA_STATUS_ERROR_ELS_REJECT</b>	The RPS ELS was rejected by the destination end port.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	The HBA referenced by <i>handle</i> does not contain an end port with Port Name <i>hbaPortWWN</i> .
<b>HBA_STATUS_ERROR</b>	Returned to indicate any problem with no required value.

## Related Information

“HBA\_GetEventBuffer Subroutine” on page 491, “HBA\_GetFC4Statistics Subroutine” on page 492, “HBA\_GetFCPStatistics Subroutine” on page 494, “HBA\_GetFcpTargetMappingV2 Subroutine” on page 495, “HBA\_GetPersistentBindingV2 Subroutine” on page 498, “HBA\_OpenAdapterByWWN Subroutine” on page 503, “HBA\_ScsiInquiryV2 Subroutine” on page 505, “HBA\_ScsiReadCapacityV2 Subroutine” on page 507, “HBA\_ScsiReportLunsV2 Subroutine” on page 509, “HBA\_SendCTPassThruV2 Subroutine” on page 511, “HBA\_SendRLS Subroutine” on page 515, “HBA\_SendRNIDV2 Subroutine” on page 517, “HBA\_SendRPL Subroutine” on page 519

---

## HBA\_SendScsiInquiry Subroutine

### Purpose

Sends a SCSI device inquiry command to a remote port of the end device.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>

HBA_STATUS HBA_SendScsiInquiry (handle, PortWWN, fclun, EVPD, PageCode, pRspBuffer, RspBufferSize, pSenseBuffer,
SenseBufferSize)
HBA_HANDLE handle;
HBA_WWN PortWWN;
HBA_UINT64 fclun;
HBA_UINT8 EVPD;
HBA_UINT32 PageCode;
void *pRspBuffer;
HBA_UINT32 RspBufferSize;
void *pSenseBuffer;
HBA_UINT32 SenseBufferSize;
```

## Description

The **HBA\_SendScsilInquiry** subroutine sends a **SCSI INQUIRY** command through a call to **ioctl** with the **SCIOLINQU** operation as its argument. The *arg* parameter for the **SCIOLINQU** operation is the address of an **scsi\_inquiry** structure. This structure is defined in the **/usr/include/sys/scsi\_buf.h** file. The *scsi\_inquiry* parameter block allows the caller to select the SCSI and LUN IDs to be queried. If successful, the inquiry data is returned in the *pRspBuffer* parameter. Successful completion occurs if a device responds at the requested SCSI ID, but the returned inquiry data must be examined to see if the requested LUN exists.

## Parameters

<i>handle</i>	Specifies the open file descriptor obtained from a successful call to the <b>open</b> subroutine.
<i>PortWWN</i>	Specifies the world wide name or port name of the target device.
<i>fcLUN</i>	Specifies the fcLUN.
<i>EVPD</i>	Specifies the value for the EVPD bit. If the value is 1, the Vital Product Data page code will be specified by the <i>PageCode</i> parameter.
<i>PageCode</i>	Specifies the Vital Product Data that is to be requested if the EVPD parameter is set to 1.
<i>pRspBuffer</i>	Points to a buffer containing the requested instruction for a send/read capacity request to an open adapter. The size of this buffer must not be greater than 255 bytes.
<i>RspBufferSize</i>	Specifies the size of the buffer to the <i>pRspBuffer</i> parameter.
<i>pSenseBuffer</i>	Points to a buffer containing the data returned from a send/read capacity request to an open adapter.
<i>SenseBufferSize</i>	Specifies the size of the buffer to the <i>pSenseBuffer</i> parameter.

## Return Values

Upon successful completion, the **HBA\_SendScsilInquiry** subroutine returns a buffer in bytes containing the SCSI inquiry, a buffer containing the SCSI sense data, and a value of **HBA\_STATUS\_OK**, or a value of 0.

If unsuccessful, an empty buffer for the SCSI inquiry, a response buffer containing the failure, and a value of **HBA\_STATUS\_ERROR**, or a value of 1 is returned.

## Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.
<b>HBA_STATUS_ERROR_ARG</b>	A value of 4 if there was a bad argument.
<b>HBA_STATUS_ERROR_ILLEGAL_WWN</b>	A value of 5 if the world wide name was not recognized.
<b>HBA_STATUS_SCSI_CHECK_CONDITION</b>	A value of 9 if a SCSI Check Condition was reported.

## Related Information

“HBA\_SendReportLUNs Subroutine” on page 513, “HBA\_SendReadCapacity Subroutine” on page 512, “HBA\_SendCTPassThru Subroutine” on page 510, “HBA\_SendRNID Subroutine” on page 516, “HBA\_SetRNIDMgmtInfo Subroutine” on page 523, and “HBA\_GetRNIDMgmtInfo Subroutine” on page 500.

SCSI Adapter Device Driver in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 2*.

Special Files in *AIX 5L Version 5.3 Files Reference*.

SCSI Subsystem Overview, A Typical Initiator-Mode SCSI Driver Transaction Sequence, Required SCSI Adapter Device Driver ioctl Commands, Understanding the Execution of Initiator I/O Requests, SCSI Error Recovery, and Understanding the `sc_buf` Structure in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

---

## HBA\_SetRNIDMgmtInfo Subroutine

### Purpose

Sends a **SCSI SET RNID** command to a remote port of the end device.

### Library

Common Host Bus Adapter Library (**libHBAAPI.a**)

### Syntax

```
#include <sys/hbaapi.h>
```

```
HBA_STATUS HBA_SetRNIDMgmtInfo (handle, info)
HBA_HANDLE handle;
HBA_MGMTINFO info;
```

### Description

The **HBA\_SetRNIDMgmtInfo** subroutine sends a **SCSI SET RNID** (Request Node Identification Data) command with the **SCIOCHBA** operation as its argument. This is done through a call to **ioctl**. The *arg* parameter for the **SCIOCHBA** operation is the address of a **scsi\_chba** structure. This structure is defined in the `/usr/include/sys/scsi_buf.h` file. The *scsi\_chba* parameter block allows the caller to select the **SET RNID** command to be sent to the adapter. The **info** structure stores the RNID data to be set. The **info** structure is defined in the `/usr/include/sys/hbaapi.h` file. The structure includes:

- `wwn`
- `unittyp`
- `PortId`
- `NumberOfAttachedNodes`
- `IPVersion`
- `UDPort`
- `IPAddress`
- `reserved`
- `TopologyDiscoveryFlags`

If successful, the SET RNID data in **info** is sent to the adapter.

### Parameters

*handle* Specifies the open file descriptor obtained from a successful call to the **open** subroutine.  
*info* Specifies the structure containing the information to be set or received from the **RNID** command

### Return Values

Upon successful completion, the **HBA\_SetRNIDMgmtInfo** subroutine returns a value of `HBA_STATUS_OK`, or a value of 0. If unsuccessful, a value of `HBA_STATUS_ERROR`, or a 1 is returned.

### Error Codes

The Storage Area Network Host Bus Adapter API subroutines return the following error codes:

<b>HBA_STATUS_OK</b>	A value of 0 on successful completion.
<b>HBA_STATUS_ERROR</b>	A value of 1 if an error occurred.
<b>HBA_STATUS_ERROR_INVALID_HANDLE</b>	A value of 3 if there was an invalid file handle.

## Related Information

“HBA\_SendScsiInquiry Subroutine” on page 521, “HBA\_SendReadCapacity Subroutine” on page 512, “HBA\_SendCTPassThru Subroutine” on page 510, “HBA\_SendReportLUNs Subroutine” on page 513, “HBA\_SendRNID Subroutine” on page 516, and “HBA\_GetRNIDMgmtInfo Subroutine” on page 500.

SCSI Adapter Device Driver in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems Volume 2*.

Special Files in *AIX 5L Version 5.3 Files Reference*.

SCSI Subsystem Overview, A Typical Initiator-Mode SCSI Driver Transaction Sequence, Required SCSI Adapter Device Driver ioctl Commands, Understanding the Execution of Initiator I/O Requests, SCSI Error Recovery, and Understanding the `sc_buf` Structure in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

---

## hpmInit, f\_hpminit, hpmStart, f\_hpmstart, hpmStop, f\_hpmstop, hpmTstart, f\_hpmtstart, hpmTstop, f\_hpmtstop, hpmGetTimeAndCounters, f\_hpmgetttimeandcounters, hpmGetCounters, f\_hpmgetcounters, hpmTerminate, and f\_hpmtterminate Subroutine

### Purpose

Provides application instrumentation for performance monitoring.

### Library

HPM Library (**libhpm.a**)

HPM Library (**libhpm.a**) includes four additional subroutines for threaded applications.

### Syntax

```
#include <libhpm.h>
```

```
void hpmInit(int taskID, char *progName);
void f_hpminit(int taskID, char *progName);
```

```
void hpmStart(int instID, char *label);
void f_hpmstart(int instID, char *label);
```

```
void hpmStop(int instID);
void f_hpmstop(int instID);
```

```
(libhpm_r only)
void hpmTstart(int instID, char *label);
void f_hpmtstart(int instID, char *label);
(libhpm_r only)
void hpmTstop(int instID);
void f_hpmtstop(int instID);
```

```
void hpmGetTimeAndCounters(int numCounters, double *time, long long *values);
void f_hpmgetttimeandcounters(int numCounters, double *time, long long *values);
```

```

void hpmGetCounters(long long *values);
void f_hpmgetcounters(long long *values);

void hpmTerminate(int taskID);
void f_hpmterminate(int taskID);

```

## Description

The **hpmInit** and **f\_hpminit** subroutines initialize tasks specified by the *taskID* and *progName* parameters.

The **hpmStart** and **f\_hpmstart** subroutines debut an instrumented code segment. If more than 100 instrumented sections are required, the **HPM\_NUM\_INST\_PTS** environment variable can be set to indicate the higher value and *instID* should be less than this value.

The **hpmStop** and **f\_hpmstop** subroutines indicate the end of the instrumented code segment *instID*. For each call to **hpmStart** and **f\_hpmstart**, there should be a corresponding call to **hpmStop** and **f\_hpmstop** with the matching *instID*.

The **hpmTstart** and **f\_hpmtstart** subroutines perform the same function as **hpmStart** and **f\_hpmstart**, but are used in threaded applications.

The **hpmTstop** and **f\_hpmtstop** subroutines perform the same function as **hpmStop** and **f\_hpmstop**, but are used in threaded applications.

The **hpmGetTimeAndCounters** and **f\_hpmgetttimeandcounters** subroutines are used to return the time in seconds and the accumulated counts since the call to **hpmInit** or **f\_hpminit**.

The **hpmGetCounters** and **f\_hpmgetcounters** subroutines return all the accumulated counts since the call to **hpmInit** or **f\_hpminit**. To minimize intrusion and overhead, the **hpmGetCounters** and **f\_hpmgetcounters** subroutines do not perform any check on the size of the *values* array. The number of counters can be obtained from the **pm\_info2\_t.maxpmcs** structure element supplied by **pm\_initialize** or by using the **pmIist -s** command. Alternatively, the application can use the current maximum value of 8.

The **hpmTerminate** and **f\_hpmterminate** subroutines end the *taskID* and generate the output. Applications that do not call **hpmTerminate** or **f\_hpmterminate**, do not generate performance information.

A summary report for each task is written by default in the *progName\_pid\_taskID.hpm* file, where *progName* is the second parameter to the **hpmInit** subroutine. If *progName* contains a space or tab character, or is otherwise invalid, a diagnostic message is written to **stderr** and the library exits with an error to avoid further problems.

The output file name can be defined with the **HPM\_OUTPUT\_NAME** environment flag. The **libhpm** still adds the file name suffix *\_taskID.hpm* for the performance files. By using this environment variable, you can specify the date and time for the output file name. For example:

```

MYDATE=$(date +"m%d:10/2/06M%S")
export HPM_OUTPUT_NAME=myprogram_$MYDATE

```

where the output file for task 27 will have the following name:

```

myprogram_yyyymmdd:HHMMSS_0027.hpm

```

The GUI and **.viz** output is deactivated by default. The aligned set of performance files named *progName\_pid\_taskID.viz* or **HPM\_OUTPUT\_NAME\_taskID.viz** will not be generated (the generation of the **.viz** file was previously activated by default and avoided with the **HPM\_VIZ\_OUTPUT = FALSE** environment variable).

## Parameters

<i>instID</i>	Specifies the instrumented section ID as an integer value greater than 0 and less than 100.
<i>label</i>	Specifies a label with a character string.
<i>numCounters</i>	Specifies an integer value that indicates the number of counters to be accessed.
<i>progName</i>	Specifies a program name using a character string label.
<i>taskID</i>	Specifies a node ID with an integer value.
<i>time</i>	Specifies a double precision float.
<i>values</i>	Specifies an array of type <b>long long</b> of size <i>numCounters</i> .

## Execution Environment

Functionality provided by the **libhpm** library is dependent upon corresponding functions in the **libpmap** and **libm** libraries. Therefore, the **-lpmap -lm** link flags must be specified when compiling applications.

## Return Values

No return values are defined.

## Error Codes

Upon failure, these **libhpm** subroutines either write error messages explicitly to **stderr** or use the PMAPI **pm\_error** function. The **pm\_error** function is called following an error return from any of the following subroutines:

- **pm\_init\_private**
- **pm\_set\_program\_mygroup**
- **pm\_stop\_mygroup**
- **pm\_get\_data\_mygroup**
- **pm\_start\_mygroup**
- **pm\_stop\_mythread**
- **pm\_get\_data\_mythread**
- **pm\_start\_mythread**
- **pm\_get\_data\_mythread**

Diagnostic messages are explicitly written to **stderr** or **stdout** in the following situations:

- **pm\_cycles** or **gettimeofday** returns an error
- The value of the *instID* parameter is invalid
- An event set is out of range
- The **libHPMevents** file or **HPM\_flags.env** file has an incorrect format
- There are internal errors.

Error messages that are not fatal are written to **stdout** or **stderr** with the text WARNING.

## Related Information

The “getrusage, getrusage64, times, or vtimes Subroutine” on page 428, “pm\_initialize Subroutine” on page 1085.

Performance Monitor API Programming in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## hsearch, hcreate, or hdestroy Subroutine

### Purpose

Manages hash tables.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <search.h>
```

```
ENTRY *hsearch ( Item, Action)
```

```
ENTRY Item;
```

```
Action Action;
```

```
int hcreate ( NumberOfElements)
```

```
size_t NumberOfElements;
```

```
void hdestroy ( )
```

### Description

**Attention:** Do not use the **hsearch**, **hcreate**, or **hdestroy** subroutine in a multithreaded environment.

The **hsearch** subroutine searches a hash table. It returns a pointer into a hash table that indicates the location of the given item. The **hsearch** subroutine uses open addressing with a multiplicative hash function.

The **hcreate** subroutine allocates sufficient space for the table. You must call the **hcreate** subroutine before calling the **hsearch** subroutine. The *NumberOfElements* parameter is an estimate of the maximum number of entries that the table will contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.

The **hdestroy** subroutine deletes the hash table. This action allows you to start a new hash table since only one table can be active at a time. After the call to the **hdestroy** subroutine, the data can no longer be considered accessible.

### Parameters

*Item*

Identifies a structure of the type **ENTRY** as defined in the **search.h** file. It contains two pointers:

**Item.key**

Points to the comparison key. The key field is of the **char** type.

**Item.data**

Points to any other data associated with that key. The data field is of the **void** type.

Pointers to data types other than the **char** type should be declared to pointer-to-character.

*Action*

Specifies the value of the *Action* enumeration parameter that indicates what is to be done with an entry if it cannot be found in the table. Values are:

**ENTER** Enters the value of the *Item* parameter into the table at the appropriate point. If the table is full, the **hsearch** subroutine returns a null pointer.

**FIND** Does not enter the value of the *Item* parameter into the table. If the value of the *Item* parameter cannot be found, the **hsearch** subroutine returns a null pointer. If the value of the *Item* parameter is found, the subroutine returns the address of the item in the hash table.

*NumberOfElements*

Provides an estimate of the maximum number of entries that the table contains. Under some circumstances, the **hcreate** subroutine may actually make the table larger than specified.

## Return Values

The **hcreate** subroutine returns a value of 0 if it cannot allocate sufficient space for the table.

## Related Information

The **bsearch** (“bsearch Subroutine” on page 123) subroutine, **lsearch** (“lsearch or lfind Subroutine” on page 764) subroutine, **malloc** (“malloc, free, realloc, calloc, mallopt, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778) subroutine, **strcmp** subroutine, **tsearch** subroutine.

Searching and Sorting Example Program and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## hypot, hypotf, or hypotl Subroutine

### Purpose

Computes the Euclidean distance function and complex absolute value.

### Libraries

IEEE Math Library (**libm.a**)

System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
```

```
double hypot ( x, y)
```

```
double x, y;
```

```
float hypotf (x, y)
```

```
float x;
```

```
float y;
```

```
long double hypotl (x, y)
```

```
long double x;
```

```
long double y;
```

### Description

The **hypot**, **hypotf** and **hypotl** subroutines compute the value of the square root of  $x^2 + y^2$  without undue overflow or underflow.



An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x*                    Specifies some double-precision floating-point value.  
*y*                    Specifies some double-precision floating-point value.

## Return Values

Upon successful completion, the **hypot**, **hypotf** and **hypotl** subroutines return the length of the hypotenuse of a right-angled triangle with sides of length *x* and *y*.

If the correct value would cause overflow, a range error occurs and the **hypotf** and **hypotl** subroutines return the value of the macro **HUGE\_VALF** and **HUGE\_VALL**, respectively.

If *x* or *y* is  $\pm\text{Inf}$ ,  $+\text{Inf}$  is returned (even if one of *x* or *y* is NaN).

If *x* or *y* is NaN, and the other is not  $\pm\text{Inf}$ , a NaN is returned.

If both arguments are subnormal and the correct result is subnormal, a range error may occur and the correct result is returned.

## Error Codes

When using the **libm.a (-lm)** library, if the correct value overflows, the **hypot** subroutine returns a **HUGE\_VAL** value.

**Note:** (**hypot (INF, value)** and **hypot (value, INF)** are both equal to **+INF** for all values, even if *value* = NaN.

When using **libmsaa.a (-lmsaa)**, if the correct value overflows, the **hypot** subroutine returns **HUGE\_VAL** and sets the global variable **errno** to **ERANGE**.

These error-handling procedures may be changed with the **matherr** subroutine when using the **libmsaa.a (-lmsaa)** library.

## Related Information

“**feclearexcept** Subroutine” on page 265, “**fetestexcept** Subroutine” on page 273, and “**class**, **\_class**, **finite**, **isnan**, or **unordered** Subroutines” on page 167.

The **matherr** (“**matherr** Subroutine” on page 789) subroutine, **sqrt** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## iconv Subroutine

### Purpose

Converts a string of characters in one character code set to another character code set.

## Library

The **iconv** Library (**libiconv.a**)

## Syntax

```
#include <iconv.h>
size_t iconv (CD, InBuf, InBytesLeft, OutBuf, OutBytesLeft)
iconv_t CD;
char **OutBuf, **InBuf;
size_t *OutBytesLeft, *InBytesLeft;
```

## Description

The **iconv** subroutine converts the string specified by the *InBuf* parameter into a different code set and returns the results in the *OutBuf* parameter. The required conversion method is identified by the *CD* parameter, which must be valid conversion descriptor returned by a previous, successful call to the **iconv\_open** subroutine.

On calling, the *InBytesLeft* parameter indicates the number of bytes in the *InBuf* buffer to be converted, and the *OutBytesLeft* parameter indicates the number of bytes remaining in the *OutBuf* buffer that do not contain converted data. These values are updated upon return so they indicate the new state of their associated buffers.

For state-dependent encodings, calling the **iconv** subroutine with the *InBuf* buffer set to null will reset the conversion descriptor in the *CD* parameter to its initial state. Subsequent calls with the *InBuf* buffer, specifying other than a null pointer, may cause the internal state of the subroutine to be altered a necessary.

## Parameters

<i>CD</i>	Specifies the conversion descriptor that points to the correct code set converter.
<i>InBuf</i>	Points to a buffer that contains the number of bytes in the <i>InBytesLeft</i> parameter to be converted.
<i>InBytesLeft</i>	Points to an integer that contains the number of bytes in the <i>InBuf</i> parameter.
<i>OutBuf</i>	Points to a buffer that contains the number of bytes in the <i>OutBytesLeft</i> parameter that has been converted.
<i>OutBytesLeft</i>	Points to an integer that contains the number of bytes in the <i>OutBuf</i> parameter.

## Return Values

Upon successful conversion of all the characters in the *InBuf* buffer and after placing the converted characters in the *OutBuf* buffer, the **iconv** subroutine returns 0, updates the *InBytesLeft* and *OutBytesLeft* parameters, and increments the *InBuf* and *OutBuf* pointers. Otherwise, it updates the variables pointed to by the parameters to indicate the extent to the conversion, returns the number of bytes still left to be converted in the input buffer, and sets the **errno** global variable to indicate the error.

## Error Codes

If the **iconv** subroutine is unsuccessful, it updates the variables to reflect the extent of the conversion before it stopped and sets the **errno** global variable to one of the following values:

**EILSEQ** Indicates an unusable character. If an input character does not belong to the input code set, no conversion is attempted on the unusable on the character. In *InBytesLeft* parameters indicates the bytes left to be converted, including the first byte of the unusable character. *InBuf* parameter points to the first byte of the unusable character sequence.

The values of *OutBuf* and *OutBytesLeft* are updated according to the number of bytes available in the output buffer that do not contain converted data.

- E2BIG** Indicates an output buffer overflow. If the *OutBuf* buffer is too small to contain all the converted characters, the character that causes the overflow is not converted. The *InBytesLeft* parameter indicates the bytes left to be converted (including the character that caused the overflow). The *InBuf* parameter points to the first byte of the characters left to convert.
- EINVAL** Indicates the input buffer was truncated. If the original value of *InBytesLeft* is exhausted in the middle of a character conversion or shift/lock block, the *InBytesLeft* parameter indicates the number of bytes undefined in the character being converted.

If an input character of shift sequence is truncated by the *InBuf* buffer, no conversion is attempted on the truncated data, and the *InBytesLeft* parameter indicates the bytes left to be converted. The *InBuf* parameter points to the first bytes if the truncated sequence. The *OutBuf* and *OutBytesLeft* values are updated according to the number of characters that were previously converted. Because some encoding may have ambiguous data, the **EINVAL** return value has a special meaning at the end of stream conversion. As such, if a user detects an EOF character on a stream that is being converted and the last return code from the **iconv** subroutine was **EINVAL**, the **iconv** subroutine should be called again, with the same *InBytesLeft* parameter and the same character string pointed to by the *InBuf* parameter as when the **EINVAL** return occurred. As a result, the converter will either convert the string as is or declare it an unusable sequence (**EILSEQ**).

## Files

`/usr/lib/nls/loc/iconv/*`

Contains code set converter methods.

## Related Information

The **iconv** command, **genxlt** command.

The **iconv\_close** (“*iconv\_close* Subroutine”) subroutine, **iconv\_open** (“*iconv\_open* Subroutine” on page 532) subroutine.

---

## iconv\_close Subroutine

### Purpose

Closes a specified code set converter.

### Library

iconv Library (**libiconv.a**)

### Syntax

```
#include <iconv.h>
```

```
int iconv_close ( CD )
iconv_t CD;
```

### Description

The **iconv\_close** subroutine closes a specified code set converter and deallocates any resources used by the converter.

### Parameters

*CD* Specifies the conversion descriptor to be closed.

## Return Values

When successful, the **iconv\_close** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **errno** global variable to indicate the error.

## Error Codes

The following error code is defined for the **iconv\_close** subroutine:

**EBADF**                      The conversion descriptor is not valid.

## Related Information

The **iconv** (“iconv Subroutine” on page 529) subroutine, **iconv\_open** (“iconv\_open Subroutine”) subroutine.

The **genxlt** command, **iconv** command.

National Language Support Overview and Converters Overview for Programming in *AIX 5L Version 5.3 National Language Support Guide and Reference*

---

## iconv\_open Subroutine

### Purpose

Opens a character code set converter.

### Library

iconv Library (**libiconv.a**)

### Syntax

```
#include <iconv.h>
```

```
iconv_t iconv_open ( ToCode, FromCode)  
const char *ToCode, *FromCode;
```

### Description

The **iconv\_open** subroutine initializes a code set converter. The code set converter is used by the **iconv** subroutine to convert characters from one code set to another. The **iconv\_open** subroutine finds the converter that performs the character code set conversion specified by the *FromCode* and *ToCode* parameters, initializes that converter, and returns a conversion descriptor of type **iconv\_t** to identify the code set converter.

The **iconv\_open** subroutine first searches the **LOCPATH** environment variable for a converter, using the two user-provided code set names, based on the file name convention that follows:

```
FromCode: "IBM-850"  
ToCode: "ISO8859-1"  
conversion file: "IBM-850_ISO8859-1"
```

The conversion file name is formed by concatenating the *ToCode* code set name onto the *FromCode* code set name, with an \_ (underscore) between them.

The **LOCPATH** environment variable contains a list of colon-separated directory names. The system default for the **LOCPATH** environment variable is:

```
LOCPATH=/usr/lib/nls/loc
```

See Locales in *AIX 5L Version 5.3 National Language Support Guide and Reference* for more information on the **LOCPATH** environment variable.

The **iconv\_open** subroutine first attempts to find the specified converter in an **iconv** subdirectory under any of the directories specified by the **LOCPATH** environment variable, for example, **/usr/lib/nls/loc/iconv**. If the **iconv\_open** subroutine cannot find a converter in any of these directories, it looks for a conversion table in an **iconvTable** subdirectory under any of the directories specified by the **LOCPATH** environment variable, for example, **/usr/lib/nls/loc/iconvTable**.

If the **iconv\_open** subroutine cannot find the specified converter in either of these locations, it returns (**iconv\_t**) -1 to the calling process and sets the **errno** global variable.

The **iconvTable** directories are expected to contain conversion tables that are the output of the **genxlt** command. The conversion tables are limited to single-byte stateless code sets. See the "List of PC, ISO, and EBCDIC Code Set Converters" in *AIX 5L Version 5.3 National Language Support Guide and Reference* for more information.

If the named converter is found, the **iconv\_open** subroutine will perform the **load** subroutine operation and initialize the converter. A converter descriptor (**iconv\_t**) is returned.

**Note:** When a process calls the **exec** subroutine or a **fork** subroutine, all of the opened converters are discarded.

The **iconv\_open** subroutine links the converter function using the **load** subroutine, which is similar to the **exec** subroutine and effectively performs a run-time linking of the converter program. Since the **iconv\_open** subroutine is called as a library function, it must ensure that security is preserved for certain programs. Thus, when the **iconv\_open** subroutine is called from a set root ID program (a program with permission **—s—s—x**), it will ignore the **LOCPATH** environment variable and search for converters only in the **/usr/lib/nls/loc/iconv** directory.

## Parameters

*ToCode* Specifies the destination code set.  
*FromCode* Specifies the originating code set.

## Return Values

A conversion descriptor (**iconv\_t**) is returned if successful. Otherwise, the subroutine returns -1, and the **errno** global variable is set to indicate the error.

## Error Codes

<b>EINVAL</b>	The conversion specified by the <i>FromCode</i> and <i>ToCode</i> parameters is not supported by the implementation.
<b>EMFILE</b>	The number of file descriptors specified by the <b>OPEN_MAX</b> configuration variable is currently open in the calling process.
<b>ENFILE</b>	Too many files are currently open in the system.
<b>ENOMEM</b>	Insufficient storage space is available.

## Files

<b>/usr/lib/nls/loc/iconv</b>	Contains loadable method converters.
<b>/usr/lib/nls/loc/iconvTable</b>	Contains conversion tables for single-byte stateless code sets.

## Related Information

The “iconv Subroutine” on page 529, “iconv\_close Subroutine” on page 531.

The **genxlt** command, **iconv** command.

Code Sets for National Language Support, List of PC, ISO, and EBCDIC Code Set Converters, National Language Support Overview , and Converters Overview for Programming in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## ilogbf, ilogbl, or ilogb Subroutine

### Purpose

Returns an unbiased exponent.

### Syntax

```
#include <math.h>
```

```
int ilogbf (x)  
float x;
```

```
int ilogbl (x)  
long double x;
```

```
int ilogb (x)  
double x;
```

### Description

The **ilogbf**, **ilogbl**, and **ilogb** subroutines return the exponent part of the *x* parameter. The return value is the integral part of  $\log_r |x|$  as a signed integral value, for nonzero *x*, where *r* is the radix of the machine’s floating-point arithmetic (*r*=2).

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Specifies the value to be computed.

### Return Values

Upon successful completion, the **ilogbf**, **ilogbl**, and **ilogb** subroutines return the exponent part of *x* as a signed integer value. They are equivalent to calling the corresponding **logb** function and casting the returned value to type **int**.

If *x* is 0, a domain error occurs, and the value **FP\_ILOGB0** is returned.

If *x* is  $\pm\text{Inf}$ , a domain error occurs, and the value **{INT\_MAX}** is returned.

If *x* is a NaN, a domain error occurs, and the value **FP\_ILOGBNAN** is returned.

If the correct value is greater than **{INT\_MAX}**, **{INT\_MAX}** is returned and a domain error occurs.

If the correct value is less than `{INT_MIN}`, `{INT_MIN}` is returned and a domain error occurs.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

`math.h` in *AIX 5L Version 5.3 Files Reference*.

---

## imaxabs Subroutine

### Purpose

Returns absolute value.

### Syntax

```
#include <inttypes.h>
```

```
intmax_t imaxabs (j)  
intmax_t j;
```

### Description

The `imaxabs` subroutine computes the absolute value of an integer *j*. If the result cannot be represented, the behavior is undefined.

### Parameters

*j* Specifies the value to be computed.

### Return Values

The `imaxabs` subroutine returns the absolute value.

## Related Information

The “`imaxdiv` Subroutine.”

`inttypes.h` File in *AIX 5L Version 5.3 Files Reference*.

---

## imaxdiv Subroutine

### Purpose

Returns quotient and remainder.

### Syntax

```
#include <inttypes.h>
```

```
imaxdiv_t imaxdiv (numer, denom)  
intmax_t numer;  
intmax_t denom;
```

### Description

The `imaxdiv` subroutine computes *numer* / *denom* and *numer* % *denom* in a single operation.

## Parameters

<i>numer</i>	Specifies the numerator value to be computed.
<i>denom</i>	Specifies the denominator value to be computed.

## Return Values

The **imaxdiv** subroutine returns a structure of type **imaxdiv\_t**, comprising both the quotient and the remainder. The structure contains (in either order) the members *quot* (the quotient) and *rem* (the remainder), each of which has type **intmax\_t**.

If either part of the result cannot be represented, the behavior is undefined.

## Related Information

The “imaxabs Subroutine” on page 535.

inttypes.h File in *AIX 5L Version 5.3 Files Reference*.

---

## IMAIXMapping Subroutine

### Purpose

Translates a pair of *Key* and *State* parameters to a string and returns a pointer to this string.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
caddr_t IMAIXMapping(IMMap, Key, State, NBytes)  
IMMap IMMap;  
KeySym Key;  
uint State;  
int * NBytes;
```

### Description

The **IMAIXMapping** subroutine translates a pair of *Key* and *State* parameters to a string and returns a pointer to this string.

This function handles the diacritic character sequence and Alt-NumPad key sequence.

## Parameters

<i>IMMap</i>	Identifies the keymap.
<i>Key</i>	Specifies the key symbol to which the string is mapped.
<i>State</i>	Specifies the state to which the string is mapped.
<i>NBytes</i>	Returns the length of the returning string.

## Return Values

If the length set by the *NBytes* parameter has a positive value, the **IMAIXMapping** subroutine returns a pointer to the returning string.

**Note:** The returning string is not null-terminated.



---

## IMAuxCreate Callback Subroutine

### Purpose

Tells the application program to create an auxiliary area.

### Syntax

```
int IMAuxCreate( IM, AuxiliaryID, UData)
IMObject IM;
caddr_t *AuxiliaryID;
caddr_t UData;
```

### Description

The **IMAuxCreate** subroutine is invoked by the input method of the operating system to create an auxiliary area. The auxiliary area can contain several different forms of data and is not restricted by the interface.

Most input methods display one auxiliary area at a time, but callbacks must be capable of handling multiple auxiliary areas.

This subroutine is provided by applications that use input methods.

### Parameters

<i>IM</i>	Indicates the input method instance.
<i>AuxiliaryID</i>	Identifies the newly created auxiliary area.
<i>UData</i>	Identifies an argument passed by the <b>IMCreate</b> subroutine.

### Return Values

On successful return of the **IMAuxCreate** subroutine, a newly created auxiliary area is set to the *AuxiliaryID* value and the **IMError** global variable is returned. Otherwise, the **IMNoError** value is returned.

### Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods, National Language Support Overview, and Using Callbacksin *AIX 5L Version 5.3 National Language Support Guide and Reference*

---

## IMAuxDestroy Callback Subroutine

### Purpose

Tells the application to destroy the auxiliary area.

### Syntax

```
int IMAuxDestroy( IM, AuxiliaryID, UData)
IMObject IM;
caddr_t AuxiliaryID;
caddr_t UData;
```

### Description

The **IMAuxDestroy** subroutine is called by the input method of the operating system to tell the application to destroy an auxiliary area.

This subroutine is provided by applications that use input methods.

## Parameters

<i>IM</i>	Indicates the input method instance.
<i>AuxiliaryID</i>	Identifies the auxiliary area to be destroyed.
<i>UData</i>	An argument passed by the <b>IMCreate</b> subroutine.

## Return Values

If an error occurs, the **IMAuxDestroy** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods, and National Language Support Overview, and Using Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMAuxDraw Callback Subroutine

### Purpose

Tells the application program to draw the auxiliary area.

### Syntax

```
int IMAuxDraw(IM, AuxiliaryID, AuxiliaryInformation, UData)
IMObject IM;
caddr_t AuxiliaryID;
IMAuxInfo * AuxiliaryInformation;
caddr_t UData;
```

### Description

The **IMAuxDraw** subroutine is invoked by the input method to draw an auxiliary area. The auxiliary area should have been previously created.

This subroutine is provided by applications that use input methods.

## Parameters

<i>IM</i>	Indicates the input method instance.
<i>AuxiliaryID</i>	Identifies the auxiliary area.
<i>AuxiliaryInformation</i>	Points to the <b>IMAuxInfo</b> structure.
<i>UData</i>	An argument passed by the <b>IMCreate</b> subroutine.

## Return Values

If an error occurs, the **IMAuxDraw** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

## Related Information

The **IMAuxCreate** (“IMAuxCreate Callback Subroutine” on page 537) subroutine, **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

## IMAuxHide Callback Subroutine

### Purpose

Tells the application program to hide an auxiliary area.

### Syntax

```
int IMAuxHide( IM, AuxiliaryID, UData)
```

```
IMObject IM;  
caddr_t AuxiliaryID;  
caddr_t UData;
```

### Description

The **IMAuxHide** subroutine is called by the input method to hide an auxiliary area.

This subroutine is provided by applications that use input methods.

### Parameters

<i>IM</i>	Indicates the input method instance.
<i>AuxiliaryID</i>	Identifies the auxiliary area to be hidden.
<i>UData</i>	An argument passed by the <b>IMCreate</b> subroutine.

### Return Values

If an error occurs, the **IMAuxHide** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

### Related Information

The **IMAuxCreate** (“IMAuxCreate Callback Subroutine” on page 537) subroutine, **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods, National Language Support Overview, and Using Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMBeep Callback Subroutine

### Purpose

Tells the application program to emit a beep sound.

### Syntax

```
int IMBeep( IM, Percent, UData)
```

```
IMObject IM;  
int Percent;  
caddr_t UData;
```

## Description

The **IMBeep** subroutine tells the application program to emit a beep sound.

This subroutine is provided by applications that use input methods.

## Parameters

<i>IM</i>	Indicates the input method instance.
<i>Percent</i>	Specifies the beep level. The value range is from -100 to 100, inclusively. A -100 value means no beep.
<i>UData</i>	An argument passed by the <b>IMCreate</b> subroutine.

## Return Values

If an error occurs, the **IMBeep** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods, National Language Support Overview, and Using Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMClose Subroutine

### Purpose

Closes the input method.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
void IMClose( IMfep)  
IMFep IMfep;
```

### Description

The **IMClose** subroutine closes the input method. Before the **IMClose** subroutine is called, all previously created input method instances must be destroyed with the **IMDestroy** subroutine, or memory will not be cleared.

### Parameters

<i>IMfep</i>	Specifies the input method.
--------------	-----------------------------

### Related Information

The **IMDestroy** (“IMDestroy Subroutine” on page 541) subroutine.

Input Method Overview and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMCreate Subroutine

### Purpose

Creates one instance of an **IMObject** object for a particular input method.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
IMObject IMCreate( IMfep, IMCallback, UData)  
IMFep IMfep;  
IMCallback *IMCallback;  
caddr_t UData;
```

### Description

The **IMCreate** subroutine creates one instance of a particular input method. Several input method instances can be created under one input method.

### Parameters

<i>IMfep</i>	Specifies the input method.
<i>IMCallback</i>	Specifies a pointer to the caller-supplied <b>IMCallback</b> structure.
<i>UData</i>	Optionally specifies an application's own information to the callback functions. With this information, the application can avoid external references from the callback functions. The input method does not change this parameter, but merely passes it to the callback functions. The <i>UData</i> parameter is usually a pointer to the application data structure, which contains the information about location, font ID, and so forth.

### Return Values

The **IMCreate** subroutine returns a pointer to the created input method instance of type **IMObject**. If the subroutine is unsuccessful, a null value is returned and the **imerrno** global variable is set to indicate the error.

### Related Information

The **IMDestroy** (“IMDestroy Subroutine”) subroutine, **IMFilter** (“IMFilter Subroutine” on page 542) subroutine, **IMLookupString** (“IMLookupString Subroutine” on page 549) subroutine, **IMProcess** (“IMProcess Subroutine” on page 550) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMDestroy Subroutine

### Purpose

Destroys an input method instance.

### Library

Input Method Library (**libIM.a**)

## Syntax

```
void IMDestroy( IM)  
IMObject IM;
```

## Description

The **IMDestroy** subroutine destroys an input method instance.

## Parameters

*IM* Specifies the input method instance to be destroyed.

## Related Information

The **IMClose** (“IMClose Subroutine” on page 540) subroutine, **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMFilter Subroutine

### Purpose

Determines if a keyboard event is used by the input method for internal processing.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
int IMFilter(Im, Key, State, String, Length)  
IMObject Im;  
Keysym Key;  
uint State, * Length;  
caddr_t * String;
```

### Description

The **IMFilter** subroutine is used to process a keyboard event and determine if the input method for this operating system uses this event. The return value indicates:

- The event is filtered (used by the input method) if the return value is **IMInputUsed**. Otherwise, the input method did not accept the event.
- Independent of the return value, a string may be generated by the keyboard event if pre-editing is complete.

**Note:** The buffer returned from the **IMFilter** subroutine is owned by the input method editor and can not continue between calls.

### Parameters

*Im* Specifies the input method instance.  
*Key* Specifies the keysym for the event.  
*State* Defines the state of the keysym. A value of 0 means that the keysym is not redefined.  
*String* Holds the returned string if one exists. A null value means that no composed string is ready.

*Length* Defines the length of the input string. If the string is not null, returns the length.

## Return Values

**IMInputUsed** The input method for this operating system filtered the event.  
**IMInputNotUsed** The input method for this operating system did not use the event.

## Related Information

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMFreeKeymap Subroutine

### Purpose

Frees resources allocated by the **IMInitializeKeymap** subroutine.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
void IMFreeKeymap( IMMap)  
IMMap IMMap;
```

### Description

The **IMFreeKeymap** subroutine frees resources allocated by the **IMInitializeKeymap** subroutine.

### Parameters

*IMMap* Identifies the keymap.

## Related Information

The **IMInitializeKeymap** (“IMInitializeKeymap Subroutine” on page 546) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMIndicatorDraw Callback Subroutine

### Purpose

Tells the application program to draw the indicator.

### Syntax

```
int IMIndicatorDraw( IM, IndicatorInformation, UData)  
IMObject IM;  
IMIndicatorInfo *IndicatorInformation;  
caddr_t UData;
```

## Description

The **IMIndicatorDraw** callback subroutine is called by the input method when the value of the indicator is changed. The application program then draws the indicator.

This subroutine is provided by applications that use input methods.

## Parameters

<i>IM</i>	Indicates the input method instance.
<i>IndicatorInformation</i>	Points to the <b>IMIndicatorInfo</b> structure that holds the current value of the indicator. The interpretation of this value varies among phonic languages. However, the input method provides a function to interpret this value.
<i>UData</i>	An argument passed by the <b>IMCreate</b> subroutine.

## Return Values

If an error happens, the **IMIndicatorDraw** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine, **IMIndicatorHide** (“IMIndicatorHide Callback Subroutine”) subroutine.

Input Methods, National Language Support Overview and Using Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMIndicatorHide Callback Subroutine

### Purpose

Tells the application program to hide the indicator.

### Syntax

```
int IMIndicatorHide( IM, UData)
IMObject IM;
caddr_t UData;
```

### Description

The **IMIndicatorHide** subroutine is called by the input method to tell the application program to hide the indicator.

This subroutine is provided by applications that use input methods.

### Parameters

<i>IM</i>	Indicates the input method instance.
<i>UData</i>	Specifies an argument passed by the <b>IMCreate</b> subroutine.

### Return Values

If an error occurs, the **IMIndicatorHide** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.



## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine, **IMIndicatorDraw** (“IMIndicatorDraw Callback Subroutine” on page 543) subroutine.

Input Methods, National Language Support Overview and Understanding Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMInitialize Subroutine

### Purpose

Initializes the input method for a particular language.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
IMFep IMInitialize( Name)  
char *Name;
```

### Description

The **IMInitialize** subroutine initializes an input method. The **IMCreate**, **IMFilter**, and **IMLookupString** subroutines use the input method to perform input processing of keyboard events in the form of keysym state modifiers. The **IMInitialize** subroutine finds the input method that performs the input processing specified by the *Name* parameter and returns an Input Method Front End Processor (**IMFep**) descriptor.

Before calling any of the key event-handling functions, the application must create an instance of an *IMObject* object using the **IMFep** descriptor. Each input method can produce one or more instances of *IMObject* object with the **IMCreate** subroutine.

When the **IMInitialize** subroutine is called, strings returned from the input method are encoded in the code set of the locale. Each **IMFep** description inherits the code set of the locale when the input method is initialized. The locale setting does not change the code set of the **IMFep** description after it is created.

The **IMInitialize** subroutine calls the **load** subroutine to load a file whose name is in the form *Name.im*. The *Name* parameter is passed to the **IMInitialize** subroutine. The loadable input method file is accessed in the directories specified by the **LOCPATH** environment variable. The default location for loadable input-method files is the **/usr/lib/nls/loc** directory. If none of the **LOCPATH** directories contain the input method specified by the *Name* parameter, the default location is searched.

**Note:** All **setuid** and **setgid** programs will ignore the **LOCPATH** environment variable.

The name of the input method file usually corresponds to the locale name, which is in the form **Language\_territory.codesest@modifier**. In the environment, the modifier is in the form **@im=modifier**. The **IMInitialize** subroutine converts the **@im=** substring to **@** when searching for loadable input-method files.

### Parameters

<i>Name</i>	Specifies the language to be used. Each input method is dynamically linked to the application program.
-------------	--

## Return Values

If **IMInitialize** succeeds, it returns an **IMFep** handle. Otherwise, null is returned and the **imerrno** global variable is set to indicate the error.

## Files

**/usr/lib/nls/loc**                      Contains loadable input-method files.

## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMInitializeKeymap Subroutine

### Purpose

Initializes the keymap associated with a specified language.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
IMMap IMInitializeKeymap( Name)
char *Name;
```

### Description

The **IMInitializeKeymap** subroutine initializes an input method keymap (imkeymap). The **IMAIXMapping** and **IMSimpleMapping** subroutines use the imkeymap to perform mapping of keysym state modifiers to strings. The **IMInitializeKeymap** subroutine finds the imkeymap that performs the keysym mapping and returns an imkeymap descriptor, **IMMap**. The strings returned by the imkeymap mapping functions are treated as unsigned bytes.

The applications that use input methods usually do not need to manage imkeymaps separately. The imkeymaps are managed internally by input methods.

The **IMInitializeKeymap** subroutine searches for an imkeymap file whose name is in the form *Name.im*. The *Name* parameter is passed to the **IMInitializeKeymap** subroutine. The imkeymap file is accessed in the directories specified by the **LOCPATH** environment variable. The default location for input method files is the **/usr/lib/nls/loc** directory. If none of the **LOCPATH** directories contain the keymap method specified by the *Name* parameter, the default location is searched.

**Note:** All **setuid** and **setgid** programs will ignore the **LOCPATH** environment variable.

The name of the imkeymap file usually corresponds to the locale name, which is in the form **Language\_territory.codesest@modifier**. In the AIXwindows environment, the modifier is in the form **@im=modifier**. The **IMInitializeKeymap** subroutine converts the **@im= substring** to **@** (at sign) when searching for loadable input method files.

## Parameters

*Name* Specifies the name of the imkeymap.

## Return Values

The **IMInitializeKeymap** subroutine returns a descriptor of type **IMMap**. Returning a null value indicates the occurrence of an error. The **IMMap** descriptor is defined in the **im.h** file as the **caddr\_t** structure. This descriptor is used for keymap manipulation functions.

## Files

**/usr/lib/nls/loc** Contains loadable input-method files.

## Related Information

The **IMFreeKeymap** (“IMFreeKeymap Subroutine” on page 543), **IMQueryLanguage** (“IMQueryLanguage Subroutine” on page 552) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMIoctl Subroutine

### Purpose

Performs a variety of control or query operations on the input method.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
int IMIoctl( IM, Operation, Argument)
IMObject IM;
int Operation;
char *Argument;
```

### Description

The **IMIoctl** subroutine performs a variety of control or query operations on the input method specified by the *IM* parameter. In addition, this subroutine can be used to control the unique function of each language input method because it provides input method-specific extensions. Each input method defines its own function.

### Parameters

*IM* Specifies the input method instance.

*Operation*  
Specifies the operation.

*Argument*  
The use of this parameter depends on which of the following operations is performed.

**IM\_Refresh**

Refreshes the text area, auxiliary areas, and indicator by calling the needed callback functions if these areas are not empty. The *Argument* parameter is not used.

**IM\_GetString**

Gets the current pre-editing string. The *Argument* parameter specifies the address of the **IMSTR** structure supplied by the caller. The callback function is invoked to clear the pre-editing if it exists.

**IM\_Clear**

Clears the text and auxiliary areas if they exist. If the *Argument* parameter is not a null value, this operation invokes the callback functions to clear the screen. The keyboard state remains the same.

**IM\_Reset**

Clears the auxiliary area if it currently exists. If the *Argument* parameter is a null value, this operation clears only the internal buffer of the input method. Otherwise, the **IMAuxHide** subroutine is called, and the input method returns to its initial state.

**IM\_ChangeLength**

Changes the maximum length of the pre-editing string.

**IM\_ChangeMode**

Sets the Processing Mode of the input method to the mode specified by the *Argument* parameter. The valid value for *Argument* is:

**IMNormalMode**

Specifies the normal mode of pre-editing.

**IMSuppressedMode**

Suppresses pre-editing.

**IM\_QueryState**

Returns the status of the text area, the auxiliary area, and the indicator. It also returns the beep status and the processing mode. The results are stored into the caller-supplied **IMQueryState** structure pointed to by the *Argument* parameter.

**IM\_QueryText**

Returns detailed information about the text area. The results are stored in the caller-supplied **IMQueryText** structure pointed to by the *Argument* parameter.

**IM\_QueryAuxiliary**

Returns detailed information about the auxiliary area. The results are stored in the caller-supplied **IMQueryAuxiliary** structure pointed to by the *Argument* parameter.

**IM\_QueryIndicator**

Returns detailed information about the indicator. The results are stored in the caller-supplied **IMQueryIndicator** structure pointed to by the *Argument* parameter.

**IM\_QueryIndicatorString**

Returns an indicator string corresponding to the current indicator. Results are stored in the caller-supplied **IMQueryIndicatorString** structure pointed to by the *Argument* parameter. The caller can request either a short or long form with the format member of the **IMQueryIndicatorString** structure.

**IM\_SupportSelection**

Informs the input method whether or not an application supports an auxiliary area selection list. The application must support selections inside the auxiliary area and determine how selections are displayed. If this operation is not performed, the input method assumes the application does not support an auxiliary area selection list.

## Return Values

The **IMloctl** subroutine returns a value to the **IMError** global variable that indicates the type of error encountered. Some error types are provided in the `/usr/include/imerrno.h` file.

## Related Information

The **IMFilter** (“IMFilter Subroutine” on page 542) subroutine, **IMLookupString** (“IMLookupString Subroutine”) subroutine, **IMProcess** (“IMProcess Subroutine” on page 550) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMLookupString Subroutine

### Purpose

Maps a *Key/State* (key symbol/state) pair to a string.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
int IMLookupString(Im, Key, State, String, Length)
IMObject Im;
KeySym Key;
uint State, * Length;
caddr_t * String;
```

### Description

The **IMLookupString** subroutine is used to map a *Key/State* pair to a localized string. It uses an internal input method keymap (**imkeymap**) file to map a keysym/modifier to a string. The string returned is encoded in the same code set as the locale of **IMObject** and IM Front End Processor.

**Note:** The buffer returned from the **IMLookupString** subroutine is owned by the input method editor and can not continue between calls.

### Parameters

<i>Im</i>	Specifies the input method instance.
<i>Key</i>	Specifies the key symbol for the event.
<i>State</i>	Defines the state for the event. A value of 0 means that the key is not redefined.
<i>String</i>	Holds the returned string, if one exists. A null value means that no composed string is ready.
<i>Length</i>	Defines the length string on input. If the string is not null, identifies the length returned.

### Return Values

<b>IMError</b>	Error encountered.
<b>IMReturnNothing</b>	No string or keysym was returned.
<b>IMReturnString</b>	String returned.

## Related Information

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMProcess Subroutine

### Purpose

Processes keyboard events and language-specific input.

### Library

Input Method Library (**libIM.a**)

**Note:** This subroutine will be removed in future releases. Use the **IMFilter** (“IMFilter Subroutine” on page 542) and **IMLookupString** (“IMLookupString Subroutine” on page 549) subroutines to process keyboard events.

### Syntax

```
int IMProcess (IM, KeySymbol, State, String, Length)
IMObject IM;
KeySym KeySymbol;
uint State;
caddr_t * String;
uint * Length;
```

### Description

This subroutine is a main entry point to the input method of the operating system. The **IMProcess** subroutine processes one keyboard event at a time. Processing proceeds as follows:

- Validates the *IM* parameter.
- Performs keyboard translation for all supported modifier states.
- Invokes internal function to do language-dependent processing.
- Performs any necessary callback functions depending on the internal state.
- Returns to application, setting the *String* and *Length* parameters appropriately.

### Parameters

<i>IM</i>	Specifies the input method instance.
<i>KeySymbol</i>	Defines the set of keyboard symbols that will be handled.
<i>State</i>	Specifies the state of the keyboard.
<i>String</i>	Holds the returned string. Returning a null value means that the input is used or discarded by the input method. <b>Note:</b> The <i>String</i> parameter is not a null-terminated string.
<i>Length</i>	Stores the length, in bytes, of the <i>String</i> parameter.

### Return Values

This subroutine returns the **IMError** global variable if an error occurs. The **IMerrno** global variable is set to indicate the error. Some of the variable values include:

<b>IMError</b>	Error occurred during this subroutine.
<b>IMTextAndAuxiliaryOff</b>	No text string in the Text area, and the Auxiliary area is not shown.
<b>IMTextOn</b>	Text string in the Text area, but no Auxiliary area.

**IMAuxiliaryOn**  
**IMTextAndAuxiliaryOn**

No text string in the Text area, and the Auxiliary area is shown.  
Text string in the Text area, and the Auxiliary is shown.

## Related Information

The **IMClose** (“IMClose Subroutine” on page 540) subroutine, **IMCreate** (“IMCreate Subroutine” on page 541) subroutine **IMFilter** (“IMFilter Subroutine” on page 542) subroutine, **IMLookupString** (“IMLookupString Subroutine” on page 549) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMProcessAuxiliary Subroutine

### Purpose

Notifies the input method of input for an auxiliary area.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
int IMProcessAuxiliary(IM, AuxiliaryID, Button, PanelRow  
    PanelColumn, ItemRow, ItemColumn, String, Length)
```

```
IMObject IM;  
caddr_t AuxiliaryID;  
uint Button;  
uint PanelRow;  
uint PanelColumn;  
uint ItemRow;  
uint ItemColumn;  
caddr_t *String;  
uint *Length;
```

### Description

The **IMProcessAuxiliary** subroutine notifies the input method instance of input for an auxiliary area.

### Parameters

<i>IM</i>	Specifies the input method instance.
<i>AuxiliaryID</i>	Identifies the auxiliary area.

<i>Button</i>	Specifies one of the following types of input:
	<b>IM_ABORT</b> Abort button is pushed.
	<b>IM_CANCEL</b> Cancel button is pushed.
	<b>IM_ENTER</b> Enter button is pushed.
	<b>IM_HELP</b> Help button is pushed.
	<b>IM_IGNORE</b> Ignore button is pushed.
	<b>IM_NO</b> No button is pushed.
	<b>IM_OK</b> OK button is pushed.
	<b>IM_RETRY</b> Retry button is pushed.
	<b>IM_SELECTED</b> Selection has been made. Only in this case do the <i>PanelRow</i> , <i>PanelColumn</i> , <i>ItemRow</i> , and <i>ItemColumn</i> parameters have meaningful values.
	<b>IM_YES</b> Yes button is pushed.
<i>PanelRow</i>	Indicates the panel on which the selection event occurred.
<i>PanelColumn</i>	Indicates the panel on which the selection event occurred.
<i>ItemRow</i>	Indicates the selected item.
<i>ItemColumn</i>	Indicates the selected item.
<i>String</i>	Holds the returned string. If a null value is returned, the input is used or discarded by the input method. Note that the <i>String</i> parameter is not a null-terminated string.
<i>Length</i>	Stores the length, in bytes, of the <i>String</i> parameter.

## Related Information

The **IMAuxCreate** (“IMAuxCreate Callback Subroutine” on page 537) subroutine.

Input Methods and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMQueryLanguage Subroutine

### Purpose

Checks to see if the specified input method is supported.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
uint IMQueryLanguage( Name)
IMLanguage Name;
```



## Description

The **IMQueryLanguage** subroutine checks to see if the input method specified by the *Name* parameter is supported.

## Parameters

*Name* Specifies the input method.

## Return Values

The **IMQueryLanguage** subroutine returns a true value if the specified input method is supported, a false value if not.

## Related Information

The **IMClose** (“IMClose Subroutine” on page 540) subroutine, **IMInitialize** (“IMInitialize Subroutine” on page 545) subroutine.

Input Methods, National Language Support Overview, Understanding Keyboard Mapping contains a list of supported languages in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMSimpleMapping Subroutine

### Purpose

Translates a pair of *KeySymbol* and *State* parameters to a string and returns a pointer to this string.

### Library

Input Method Library (**libIM.a**)

### Syntax

```
caddr_t IMSimpleMapping (IMMap, KeySymbol, State, NBytes)
IMMap  IMMap;
KeySym KeySymbol;
uint   State;
int *  NBytes;
```

### Description

Like the **IMAIXMapping** subroutine, the **IMSimpleMapping** subroutine translates a pair of *KeySymbol* and *State* parameters to a string and returns a pointer to this string. The parameters have the same meaning as those in the **IMAIXMapping** subroutine.

The **IMSimpleMapping** subroutine differs from the **IMAIXMapping** subroutine in that it does not support the diacritic character sequence or the Alt-NumPad key sequence.

### Parameters

<i>IMMap</i>	Identifies the keymap.
<i>KeySymbol</i>	Key symbol to which the string is mapped.
<i>State</i>	Specifies the state to which the string is mapped.
<i>NBytes</i>	Returns the length of the returning string.

## Related Information

The **IMAIXMapping** (“IMAIXMapping Subroutine” on page 536) subroutine, **IMFreeKeymap** (“IMFreeKeymap Subroutine” on page 543) subroutine, **IMInitializeKeymap** (“IMInitializeKeymap Subroutine” on page 546) subroutine.

Input Method Overview and National Language Support Overview for Programming in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## IMTextCursor Callback Subroutine

### Purpose

Asks the application to move the text cursor.

### Syntax

```
int IMTextCursor(IM, Direction, Cursor, UData)
IMObject IM;
uint Direction;
int * Cursor;
caddr_t UData;
```

### Description

The **IMTextCursor** subroutine is called by the Input Method when the Cursor Up or Cursor Down key is input to the **IMFilter** and **IMLookupString** subroutines.

This subroutine sets the new display cursor position in the text area to the integer pointed to by the *Cursor* parameter. The cursor position is relative to the top of the text area. A value of -1 indicates the cursor should not be moved.

Because the input method does not know the actual length of the screen it always treats a text string as one-dimensional (a single line). However, in the terminal emulator, the text string sometimes wraps to the next line. The **IMTextCursor** subroutine performs this conversion from single-line to multiline text strings. When you move the cursor up or down, the subroutine interprets the cursor position on the text string relative to the input method.

This subroutine is provided by applications that use input methods.

### Parameters

<i>IM</i>	Indicates the Input Method instance.
<i>Direction</i>	Specifies up or down.
<i>Cursor</i>	Specifies the new cursor position or -1.
<i>UData</i>	Specifies an argument passed by the <b>IMCreate</b> subroutine.

### Return Values

If an error occurs, the **IMTextCursor** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine, **IMFilter** (“IMFilter Subroutine” on page 542) subroutine, **IMLookupString** (“IMLookupString Subroutine” on page 549) subroutine, **IMTextDraw** (“IMTextDraw Callback Subroutine” on page 555) subroutine.

## IMTextDraw Callback Subroutine

### Purpose

Tells the application program to draw the text string.

### Syntax

```
int IMTextDraw( IM, TextInfo, UData)
IMObject IM;
IMTextInfo *TextInfo;
caddr_t UData;
```

### Description

The **IMTextDraw** subroutine is invoked by the Input Method whenever it needs to update the screen with its internal string. This subroutine tells the application program to draw the text string.

This subroutine is provided by applications that use input methods.

### Parameters

<i>IM</i>	Indicates the input method instance.
<i>TextInfo</i>	Points to the <b>IMTextInfo</b> structure.
<i>UData</i>	An argument passed by the <b>IMCreate</b> subroutine.

### Return Values

If an error occurs, the **IMTextDraw** subroutine returns the **IMError** global variable. Otherwise, the **IMNoError** value is returned.

### Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine.

Input Methods, National Language Support Overview, and Using Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMTextHide Callback Subroutine

### Purpose

Tells the application program to hide the text area.

### Syntax

```
int IMTextHide( IM, UData)
IMObject IM;
caddr_t UData;
```

### Description

The **IMTextHide** subroutine is called by the input method when the text area should be cleared. This subroutine tells the application program to hide the text area.

This subroutine is provided by applications that use input methods.

## Parameters

*IM* Indicates the input method instance.  
*UData* Specifies an argument passed by the **IMCreate** subroutine.

## Return Values

If an error occurs, the **IMTextHide** subroutine returns an **IMError** value. Otherwise, an **IMNoError** value is returned.

## Related Information

The **IMTextDraw** (“IMTextDraw Callback Subroutine” on page 555) subroutine.

Input Methods, National Language Support Overview, and Using Callbacks in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## IMTextStart Callback Subroutine

### Purpose

Notifies the application program of the length of the pre-editing space.

### Syntax

```
int IMTextStart( IM, Space, UData)
IMObject IM;
int *Space;
caddr_t UData;
```

### Description

The **IMTextStart** subroutine is called by the input method when the pre-editing is started, but prior to calling the **IMTextDraw** callback subroutine. This subroutine notifies the input method of the length, in terms of bytes, of pre-editing space. It sets the length of the available space ( $\geq 0$ ) on the display to the integer pointed to by the *Space* parameter. A value of -1 indicates that the pre-editing space is dynamic and has no limit.

This subroutine is provided by applications that use input methods.

## Parameters

*IM* Indicates the input method instance.  
*Space* Maximum length of pre-editing string.  
*UData* An argument passed by the **IMCreate** subroutine.

## Related Information

The **IMCreate** (“IMCreate Subroutine” on page 541) subroutine, **IMTextDraw** (“IMTextDraw Callback Subroutine” on page 555) subroutine.

Input Methods, Using Callbacks, and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## inet\_aton Subroutine

### Purpose

Converts an ASCII string into an Internet address.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
int inet_aton ( CharString, InternetAddr)
char * CharString;
struct in_addr * InternetAddr;
```

### Description

The **inet\_aton** subroutine takes an ASCII string representing the Internet address in dot notation and converts it into an Internet address.

All applications containing the **inet\_aton** subroutine must be compiled with **\_BSD** set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>CharString</i>	Contains the ASCII string to be converted to an Internet address.
<i>InternetAddr</i>	Contains the Internet address that was converted from the ASCII string.

### Return Values

Upon successful completion, the **inet\_aton** subroutine returns 1 if *CharString* is a valid ASCII representation of an Internet address.

The **inet\_aton** subroutine returns 0 if *CharString* is not a valid ASCII representation of an Internet address.

### Files

<i>/etc/hosts</i>	Contains host names.
<i>/etc/networks</i>	Contains network names.

### Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_addr** subroutine, **inet\_lnaof** subroutine, **inet\_makeaddr** subroutine, **inet\_network** subroutine, **inet\_ntoa** subroutine, **sethostent** subroutine, **setnetent** subroutine.

Sockets Overview and Network Address Translation in *AIX 5L Version 5.3 Communications Programming Concepts*.

---

## initgroups Subroutine

### Purpose

Initializes supplementary group ID.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int initgroups ( User, BaseGID)
const char *User;
int BaseGID;
```

### Description

**Attention:** The **initgroups** subroutine uses the **getgrent** and **getpwent** family of subroutines. If the program that invokes the **initgroups** subroutine uses any of these subroutines, calling the **initgroups** subroutine overwrites the static storage areas used by these subroutines.

The **initgroups** subroutine reads the defined group membership of the specified *User* parameter and sets the supplementary group ID of the current process to that value. The *BaseGID* parameter is always included in the supplementary group ID. The supplementary group is normally the principal user's group. If the user is in more than **NGROUPS\_MAX** groups, set in the **limits.h** file, only **NGROUPS\_MAX** groups are set, including the *BaseGID* group.

### Parameters

<i>User</i>	Identifies a user.
<i>BaseGID</i>	Specifies an additional group to include in the group set.

### Return Values

0	Indicates that the subroutine was success.
-1	Indicates that the subroutine failed. The <b>errno</b> global variable is set to indicate the error.

### Related Information

The **getgid** (“getgid, getegid or gegidx Subroutine” on page 370) subroutine, **getgrent**, **getgrgid**, **getgrnam**, **putgrent**, **setgrent**, or **endgrent** (“getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371) subroutine, **getgroups** (“getgroups Subroutine” on page 383) subroutine, **setgroups** subroutine.

The **groups** command, **setgroups** command.

List of Security and Auditing Subroutines, Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## initialize Subroutine

### Purpose

Performs printer initialization.

## Library

None (provided by the formatter).

## Syntax

```
#include <piostruct.h>
```

```
int initialize ()
```

## Description

The **initialize** subroutine is invoked by the formatter driver after the **setup** subroutine returns.

If the **-j** flag passed from the **qprt** command has a nonzero value (true), the **initialize** subroutine uses the **piocmdout** subroutine to send a command string to the printer. This action initializes the printer to the proper state for printing the file. Any variables referenced by the command string should be the attribute values from the database, overridden by values from the command line.

If the **-j** flag passed from the **qprt** command has a nonzero value (true), any necessary fonts should be downloaded.

## Return Values

0 Indicates a successful operation.

If the **initialize** subroutine detects an error, it uses the **piomsgout** subroutine to invoke an error message. It then invokes the **pioexit** subroutine with a value of **PIOEXITBAD**.

**Note:** If either the **piocmdout** or **piogetstr** subroutine detects an error, it issues its own error messages and terminates the print job.

## Related Information

The **piocmdout** subroutine, **pioexit** subroutine, **piogetstr** subroutine, **piomsgout** subroutine, **setup** subroutine.

Adding a New Printer Type to Your System, Printer Addition Management Subsystem: Programming Overview, Understanding Embedded References in Printer Attribute Strings in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

Print formatter example in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## insque or remque Subroutine

### Purpose

Inserts or removes an element in a queue.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <search.h>
```

```
insque ( Element, Pred)
void *Element, *Pred;
remque (Element)
void *Element;
```

## Description

The **insque** and **remque** subroutines manipulate queues built from double-linked lists. Each element in the queue must be in the form of a **qelem** structure. The **next** and **prev** elements of that structure must point to the elements in the queue immediately before and after the element to be inserted or deleted.

The **insque** subroutine inserts the element pointed to by the *Element* parameter into a queue immediately after the element pointed to by the *Pred* parameter.

The **remque** subroutine removes the element defined by the *Element* parameter from a queue.

## Parameters

*Pred*                Points to the element in the queue immediately before the element to be inserted or deleted.  
*Element*           Points to the element in the queue immediately after the element to be inserted or deleted.

## Related Information

Searching and Sorting Example Program in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## install\_lwcf\_handler Subroutine

### Purpose

Registers the signal handler to dump a lightweight core file for signals that normally cause the generation of a core file.

### Library

PTools Library (**libptools\_ptr.a**)

### Syntax

```
void install_lwcf_handler (void);
```

### Description

The **install\_lwcf\_handler** subroutine registers the signal handler to dump a lightweight core file for signals that normally cause a core file to be generated. The format of lightweight core files complies with the Parallel Tools Consortium Lightweight Core File Format.

The **install\_lwcf\_handler** subroutine uses the **LIGHTWEIGHT\_CORE** environment variable to determine the target lightweight core file. If the **LIGHTWEIGHT\_CORE** environment variable is defined, a lightweight core file will be generated. Otherwise, a normal core file will be generated.

If the **LIGHTWEIGHT\_CORE** environment variable is defined without a value, the lightweight core file is assigned the default file name **lw\_core** and is created under the current working directory if it does not already exist.



If the **LIGHTWEIGHT\_CORE** environment variable is defined with a value of **STDERR**, the lightweight core file is output to the standard error output device of the process. Keyword **STDERR** is not case-sensitive.

If the **LIGHTWEIGHT\_CORE** environment variable is defined with the value of a character string other than **STDERR**, the string is used as a path name for the lightweight core file generated.

If the target lightweight core file already exists, the traceback information is appended to the file.

The **install\_lwcf\_handler** subroutine can be called directly from an application to register the signal handler. Alternatively, linker option **-binitfni:install\_lwcf\_handler** can be used when linking an application, which specifies to execute the **install\_lwcf\_handler** subroutine when the application is initialized. The advantage of the second method is that the application code does not need to change to invoke the **install\_lwcf\_handler** subroutine.

## Related Information

The **mt\_trce** and **sigaction** subroutines.

---

## ioctl, ioctlx, ioctl32, or ioctl32x Subroutine

### Purpose

Performs control functions associated with open file descriptors.

### Library

Standard C Library (**libc.a**)

BSD Library (**libbsd.a**)

### Syntax

```
#include <sys/ioctl.h>
#include <sys/types.h>
#include <unistd.h>
#include <stropts.h>
```

```
int ioctl (FileDescriptor, Command, Argument)
int FileDescriptor, Command;
void * Argument;
```

```
int ioctlx (FileDescriptor, Command, Argument, Ext)
int FileDescriptor, Command;
void *Argument;
int Ext;
```

```
int ioctl32 (FileDescriptor, Command, Argument)
int FileDescriptor, Command;
unsigned int Argument;
```

```
int ioctl32x (FileDescriptor, Command, Argument, Ext)
int FileDescriptor, Command;
unsigned int Argument;
unsigned int Ext;
```

## Description

The **ioctl** subroutine performs a variety of control operations on the object associated with the specified open file descriptor. This function is typically used with character or block special files, sockets, or generic device support such as the **termio** general terminal interface.

The control operation provided by this function call is specific to the object being addressed, as are the data type and contents of the *Argument* parameter. The **ioctlx** form of this function can be used to pass an additional extension parameter to objects supporting it. The **ioctl132** and **ioctl132x** forms of this function behave in the same way as **ioctl** and **ioctlx**, but allow 64-bit applications to call the **ioctl** routine for an object that does not normally work with 64-bit applications.

Performing an **ioctl** function on a file descriptor associated with an ordinary file results in an error being returned.

## Parameters

<i>FileDescriptor</i>	Specifies the open file descriptor for which the control operation is to be performed.
<i>Command</i>	Specifies the control function to be performed. The value of this parameter depends on which object is specified by the <i>FileDescriptor</i> parameter.
<i>Argument</i>	Specifies additional information required by the function requested in the <i>Command</i> parameter. The data type of this parameter (a <b>void</b> pointer) is object-specific, and is typically used to point to an object device-specific data structure. However, in some device-specific instances, this parameter is used as an integer.
<i>Ext</i>	Specifies an extension parameter used with the <b>ioctlx</b> subroutine. This parameter is passed on to the object associated with the specified open file descriptor. Although normally of type <b>int</b> , this parameter can be used as a pointer to a device-specific structure for some devices.

## File Input/Output (FIO) ioctl Command Values

A number of file input/output (FIO) **ioctl** commands are available to enable the **ioctl** subroutine to function similar to the **fcntl** subroutine:

### FIOCLEX and FIONCLEX

Manipulate the **close-on-exec** flag to determine if a file descriptor should be closed as part of the normal processing of the **exec** subroutine. If the flag is set, the file descriptor is closed. If the flag is clear, the file descriptor is left open.

The following code sample illustrates the use of the **fcntl** subroutine to set and clear the **close-on-exec** flag:

```
/* set the close-on-exec flag for fd1 */
fcntl(fd1, F_SETFD, FD_CLOEXEC);
/* clear the close-on-exec flag for fd2 */
fcntl(fd2, F_SETFD, 0);
```

Although the **fcntl** subroutine is normally used to set the **close-on-exec** flag, the **ioctl** subroutine may be used if the application program is linked with the Berkeley Compatibility Library (**libbsd.a**) or the Berkeley Thread Safe Library (**libbsd\_r.a**) (4.2.1 and later versions). The following **ioctl** code fragment is equivalent to the preceding **fcntl** fragment:

```
/* set the close-on-exec flag for fd1 */
ioctl(fd1, FIOCLEX, 0);
/* clear the close-on-exec flag for fd2 */
ioctl(fd2, FIONCLEX, 0);
```

The third parameter to the **ioctl** subroutine is not used for the **FIOCLEX** and **FIONCLEX** **ioctl** commands.

**FIONBIO**

Enables nonblocking I/O. The effect is similar to setting the **O\_NONBLOCK** flag with the **fcntl** subroutine. The third parameter to the **ioctl** subroutine for this command is a pointer to an integer that indicates whether nonblocking I/O is being enabled or disabled. A value of 0 disables non-blocking I/O. Any nonzero value enables nonblocking I/O. A sample code fragment follows:

```
int flag;
/* enable NBI0 for fd1 */
flag = 1;
ioctl(fd1,FIONBIO,&flag);
/* disable NBI0 for fd2 */
flag = 0;
ioctl(fd2,FIONBIO,&flag);
```

**FIONREAD**

Determines the number of bytes that are immediately available to be read on a file descriptor. The third parameter to the **ioctl** subroutine for this command is a pointer to an integer variable where the byte count is to be returned. The following sample code illustrates the proper use of the **FIONREAD** **ioctl** command:

```
int nbytes;
ioctl(fd,FIONREAD,&nbytes);
```

**FIOASYNC**

Enables a simple form of asynchronous I/O notification. This command causes the kernel to send **SIGIO** signal to a process or a process group when I/O is possible. Only sockets, ttys, and pseudo-ttys implement this functionality.

The third parameter of the **ioctl** subroutine for this command is a pointer to an integer variable that indicates whether the asynchronous I/O notification should be enabled or disabled. A value of 0 disables I/O notification; any nonzero value enables I/O notification. A sample code segment follows:

```
int flag;
/* enable ASYNC on fd1 */
flag = 1;
ioctl(fd, FIOASYNC,&flag);
/* disable ASYNC on fd2 */
flag = 0;
ioctl(fd,FIOASYNC,&flag);
```

**FIOSETOWN**

Sets the recipient of the **SIGIO** signals when asynchronous I/O notification (**FIOASYNC**) is enabled. The third parameter to the **ioctl** subroutine for this command is a pointer to an integer that contains the recipient identifier. If the value of the integer pointed to by the third parameter is negative, the value is assumed to be a process group identifier. If the value is positive, it is assumed to be a process identifier.

Sockets support both process groups and individual process recipients, while ttys and psuedo-ttys support only process groups. Attempts to specify an individual process as the recipient will be converted to the process group to which the process belongs. The following code example illustrates how to set the recipient identifier:

```
int owner;
owner = -getpgrp();
ioctl(fd,FIOSETOWN,&owner);
```

**Note:** In this example, the asynchronous I/O signals are being enabled on a process group basis. Therefore, the value passed through the owner parameter must be a negative number.

The following code sample illustrates enabling asynchronous I/O signals to an individual process:

```
int owner;
owner = getpid();
ioctl(fd,FIOSETOWN,&owner);
```

## FIOGETOWN

Determines the current recipient of the asynchronous I/O signals of an object that has asynchronous I/O notification (**FIOASYNC**) enabled. The third parameter to the **ioctl** subroutine for this command is a pointer to an integer used to return the owner ID. For example:

```
int owner;
ioctl(fd,FIOGETOWN,&owner);
```

If the owner of the asynchronous I/O capability is a process group, the value returned in the reference parameter is negative. If the owner is an individual process, the value is positive.

## Return Values

If the **ioctl** subroutine fails, a value of -1 is returned. The **errno** global variable is set to indicate the error.

The **ioctl** subroutine fails if one or more of the following are true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid open file descriptor.
<b>EFAULT</b>	The <i>Argument</i> or <i>Ext</i> parameter is used to point to data outside of the process address space.
<b>EINTR</b>	A signal was caught during the <b>ioctl</b> or <b>ioctlx</b> subroutine and the process had not enabled re-startable subroutines for the signal.
<b>EINTR</b>	A signal was caught during the <b>ioctl</b> , <b>ioctlx</b> , <b>ioctl32</b> , or <b>ioctl32x</b> subroutine and the process had not enabled re-startable subroutines for the signal.
<b>EINVAL</b>	The <i>Command</i> or <i>Argument</i> parameter is not valid for the specified object.
<b>ENOTTY</b>	The <i>FileDescriptor</i> parameter is not associated with an object that accepts control functions.
<b>ENODEV</b>	The <i>FileDescriptor</i> parameter is associated with a valid character or block special file, but the supporting device driver does not support the <b>ioctl</b> function.
<b>ENXIO</b>	The <i>FileDescriptor</i> parameter is associated with a valid character or block special file, but the supporting device driver is not in the configured state.

Object-specific error codes are defined in the documentation for associated objects.

## Related Information

The **ddioctl** device driver entry point and the **fp\_ioctl** kernel service in *AIX 5L Version 5.3 Technical Reference: Kernel and Subsystems*.

The Special Files Overview in *AIX 5L Version 5.3 Files Reference*.

The Input and Output Handling Programmer's Overview, the tty Subsystem Overview, in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

The Sockets Overview and Understanding Socket Data Transfer in *AIX 5L Version 5.3 Communications Programming Concepts*.

---

## isblank Subroutine

### Purpose

Tests for a blank character.

## Syntax

```
#include <ctype.h>
```

```
int isblank (c)  
int c;
```

## Description

The **isblank** subroutine tests whether the *c* parameter is a character of class **blank** in the program's current locale.

The *c* parameter is a type **int**, the value of which the application shall ensure is a character representable as an **unsigned char** or equal to the value of the macro EOF. If the parameter has any other value, the behavior is undefined.

## Parameters

*c* Specifies the character to be tested.

## Return Values

The **isblank** subroutine returns nonzero if *c* is a <blank>; otherwise, it returns 0.

## Related Information

The “ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines” on page 208.

setlocale Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## isendwin Subroutine

### Purpose

Determines whether the **endwin** subroutine was called without any subsequent refresh calls.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
isendwin()
```

### Description

The **isendwin** subroutine determines whether the **endwin** subroutine was called without any subsequent refresh calls. If the **endwin** was called without any subsequent calls to the **wrefresh** or **doupdate** subroutines, the **isendwin** subroutine returns TRUE.

### Return Values

**TRUE** Indicates the **endwin** subroutine was called without any subsequent calls to the **wrefresh** or **doupdate** subroutines.  
**FALSE** Indicates subsequent calls to the refresh subroutines.

## Related Information

The **douupdate** subroutine, **endwin** subroutine, **wrefresh** subroutine.

Curses Overview for Programming, Initializing Curses, List of Curses Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## isfinite Macro

### Purpose

Tests for finite value.

### Syntax

```
#include <math.h>
```

```
int isfinite (x)  
real-floating x;
```

### Description

The **isfinite** macro determines whether its argument has a finite value (zero, subnormal, or normal, and not infinite or NaN). An argument represented in a format wider than its semantic type is converted to its semantic type. Determination is based on the type of the argument.

### Parameters

*x* Specifies the value to be tested.

### Return Values

The **isfinite** macro returns a nonzero value if its argument has a finite value.

## Related Information

“fpclassify Macro” on page 309, “isinf Subroutine” on page 568, “class, \_class, finite, isnan, or unordered Subroutines” on page 167, “isnormal Macro” on page 570.

The signbit Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## isgreater Macro

### Purpose

Tests if *x* is greater than *y*.

### Syntax

```
#include <math.h>
```

```
int isgreater (x, y)  
real-floating x;  
real-floating y;
```

## Description

The **isgreater** macro determines whether its first argument is greater than its second argument. The value of **isgreater**( $x$ ,  $y$ ) is equal to  $(x) > (y)$ ; however, unlike  $(x) > (y)$ , **isgreater**( $x$ ,  $y$ ) does not raise the invalid floating-point exception when  $x$  and  $y$  are unordered.

## Parameters

$x$  Specifies the first value to be compared.  
 $y$  Specifies the first value to be compared.

## Return Values

Upon successful completion, the **isgreater** macro returns the value of  $(x) > (y)$ .

If  $x$  or  $y$  is NaN, 0 is returned.

## Related Information

“isgreaterequal Subroutine,” “isless Macro” on page 568, “islessequal Macro” on page 569, “islessgreater Macro” on page 570, and “isunordered Macro” on page 571.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## isgreaterequal Subroutine

### Purpose

Tests if  $x$  is greater than or equal to  $y$ .

### Syntax

```
#include <math.h>

int isgreaterequal (x, y)
real-floating x;
real-floating y;
```

### Description

The **isgreaterequal** macro determines whether its first argument is greater than or equal to its second argument. The value of **isgreaterequal**( $x$ ,  $y$ ) is equal to  $(x) >= (y)$ ; however, unlike  $(x) >= (y)$ , **isgreaterequal**( $x$ ,  $y$ ) does not raise the invalid floating-point exception when  $x$  and  $y$  are unordered.

### Parameters

$x$  Specifies the first value to be compared.  
 $y$  Specifies the second value to be compared.

### Return Values

Upon successful completion, the **isgreaterequal** macro returns the value of  $(x) >= (y)$ .

If  $x$  or  $y$  is NaN, 0 is returned.

## Related Information

“isgreater Macro” on page 566, “isless Macro,” “islessequal Macro” on page 569, “islessgreater Macro” on page 570, and “isunordered Macro” on page 571.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## isinf Subroutine

### Purpose

Tests for infinity.

### Syntax

```
#include <math.h>
```

```
int isinf (x)  
real-floating x;
```

### Description

The **isinf** macro determines whether its argument value is an infinity (positive or negative). An argument represented in a format wider than its semantic type is converted to its semantic type. Determination is based on the type of the argument.

### Parameters

*x* Specifies the value to be checked.

### Return Values

The **isinf** macro returns a nonzero value if its argument has an infinite value.

## Related Information

“fpclassify Macro” on page 309, “isfinite Macro” on page 566, “class, \_class, finite, isnan, or unordered Subroutines” on page 167, “isnormal Macro” on page 570.

The signbit Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## isless Macro

### Purpose

Tests if *x* is less than *y*.

### Syntax

```
#include <math.h>  
int isless (x, y)  
real-floating x;  
real-floating y;
```



## Description

The **isless** macro determines whether its first argument is less than its second argument. The value of **isless**( $x$ ,  $y$ ) is equal to  $(x) < (y)$ ; however, unlike  $(x) < (y)$ , **isless**( $x$ ,  $y$ ) does not raise the invalid floating-point exception when  $x$  and  $y$  are unordered.

## Parameters

$x$  Specifies the first value to be compared.  
 $y$  Specifies the second value to be compared.

## Return Values

Upon successful completion, the **isless** macro returns the value of  $(x) < (y)$ .

If  $x$  or  $y$  is NaN, 0 is returned.

## Related Information

“isgreater Macro” on page 566, “isgreaterequal Subroutine” on page 567, “islessequal Macro,” “islessgreater Macro” on page 570, and “isunordered Macro” on page 571.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## islessequal Macro

### Purpose

Tests if  $x$  is less than or equal to  $y$ .

### Syntax

```
#include <math.h>

int islessequal (x, y)
real-floating x;
real-floating y;
```

### Description

The **islessequal** macro determines whether its first argument is less than or equal to its second argument. The value of **islessequal**( $x$ ,  $y$ ) is equal to  $(x) <= (y)$ ; however, unlike  $(x) <= (y)$ , **islessequal**( $x$ ,  $y$ ) does not raise the invalid floating-point exception when  $x$  and  $y$  are unordered.

### Parameters

$x$  Specifies the first value to be compared.  
 $y$  Specifies the second value to be compared.

### Return Values

Upon successful completion, the **islessequal** macro returns the value of  $(x) <= (y)$ .

If  $x$  or  $y$  is NaN, 0 is returned.

## Related Information

“isgreater Macro” on page 566, “isgreaterequal Subroutine” on page 567, “islessequal Macro” on page 569, “islessgreater Macro,” and “isunordered Macro” on page 571.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## islessgreater Macro

### Purpose

Tests if  $x$  is less than or greater than  $y$ .

### Syntax

```
#include <math.h>

int islessgreater (x, y)
real-floating x;
real-floating y;
```

### Description

The **islessgreater** macro determines whether its first argument is less than or greater than its second argument. The **islessgreater**( $x$ ,  $y$ ) macro is similar to  $(x) < (y) \parallel (x) > (y)$ ; however, **islessgreater**( $x$ ,  $y$ ) does not raise the invalid floating-point exception when  $x$  and  $y$  are unordered (nor does it evaluate  $x$  and  $y$  twice).

### Parameters

$x$  Specifies the first value to be compared.  
 $y$  Specifies the second value to be compared.

### Return Values

Upon successful completion, the **islessgreater** macro returns the value of  $(x) < (y) \parallel (x) > (y)$ .

If  $x$  or  $y$  is NaN, 0 is returned.

## Related Information

“isgreater Macro” on page 566, “isgreaterequal Subroutine” on page 567, “isless Macro” on page 568, “islessequal Macro” on page 569, and “isunordered Macro” on page 571.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## isnormal Macro

### Purpose

Tests for a normal value.

### Syntax

```
#include <math.h>

int isnormal (x)
real-floating x;
```

## Description

The **isnormal** macro determines whether its argument value is normal (neither zero, subnormal, infinite, nor NaN) or not. An argument represented in a format wider than its semantic type is converted to its semantic type. Determination is based on the type of the argument.

## Parameters

*x* Specifies the value to be tested.

## Return Values

The **isnormal** macro returns a nonzero value if its argument has a normal value.

## Related Information

“fpclassify Macro” on page 309, “isfinite Macro” on page 566, “isinf Subroutine” on page 568, “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

The signbit Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## isunordered Macro

### Purpose

Tests if arguments are unordered.

### Syntax

```
#include <math.h>
int isunordered (x, y)
real-floating x;
real-floating y;
```

### Description

The **isunordered** macro determines whether its arguments are unordered.

### Parameters

*x* Specifies the first value in the order.  
*y* Specifies the second value in the order.

### Return Values

Upon successful completion, the **isunordered** macro returns 1 if its arguments are unordered, and 0 otherwise.

If *x* or *y* is NaN, 0 is returned.

### Related Information

“isgreater Macro” on page 566, “isgreaterequal Subroutine” on page 567, “isless Macro” on page 568, “islessequal Macro” on page 569, and “islessgreater Macro” on page 570.

---

## iswalnum, iswalpha, iswcntrl, iswdigit, iswgraph, iswlower, iswprint, iswpunct, iswspace, iswupper, or iswxdigit Subroutine

### Purpose

Tests a wide character for membership in a specific character class.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wchar.h>
int iswalnum (WC)
wint_t WC;
int iswalpha (WC)
wint_t WC;
int iswcntrl (WC)
wint_t WC;
int iswdigit (WC)
wint_t WC;
int iswgraph (WC)
wint_t WC;
int iswlower (WC)
wint_t WC;
int iswprint (WC)
wint_t WC;
int iswpunct (WC)
wint_t WC;
int iswspace (WC)
wint_t WC;
int iswupper (WC)
wint_t WC;
int iswxdigit (WC)
wint_t WC;
```

### Description

The **isw** subroutines check the character class status of the wide character code specified by the *WC* parameter. Each subroutine tests to see if a wide character is part of a different character class. If the wide character is part of the character class, the **isw** subroutine returns true; otherwise, it returns false.

Each subroutine is named by adding the **isw** prefix to the name of the character class that the subroutine tests. For example, the **iswalpha** subroutine tests whether the wide character specified by the *WC* parameter is an alphabetic character. The character classes are defined as follows:

<b>alnum</b>	Alphanumeric character.
<b>alpha</b>	Alphabetic character.
<b>cntrl</b>	Control character. No characters in the <b>alpha</b> or <b>print</b> classes are included.
<b>digit</b>	Numeric digit character.
<b>graph</b>	Graphic character for printing, not including the space character or <b>cntrl</b> characters. Includes all characters in the <b>digit</b> and <b>punct</b> classes.
<b>lower</b>	Lowercase character. No characters in <b>cntrl</b> , <b>digit</b> , <b>punct</b> , or <b>space</b> are included.
<b>print</b>	Print character. All characters in the <b>graph</b> class are included, but no characters in <b>cntrl</b> are included.

<b>punct</b>	Punctuation character. No characters in the <b>alpha</b> , <b>digit</b> , or <b>cntrl</b> classes, or the space character are included.
<b>space</b>	Space characters.
<b>upper</b>	Uppercase character.
<b>xdigit</b>	Hexadecimal character.

## Parameters

*WC* Specifies a wide character for testing.

## Return Values

If the wide character tested is part of the particular character class, the **isw** subroutine returns a nonzero value; otherwise it returns a value of 0.

## Related Information

The **iswctype** subroutine, (“iswctype or is\_wctype Subroutine” on page 574)**setlocale** subroutine, **towlower** subroutine, **toupper** subroutine **wctype** subroutine.

Subroutines, Example Programs, and Libraries, Wide Character Classification Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## iswblank Subroutine

### Purpose

Tests for a blank wide-character code.

### Syntax

```
#include <wctype.h>
```

```
int iswblank (wc)
wint_t wc;
```

### Description

The **iswblank** subroutine tests whether the *wc* parameter is a wide-character code representing a character of class **blank** in the program’s current locale.

The *wc* parameter is a **wint\_t**, the value of which the application ensures is a wide-character code corresponding to a valid character in the current locale, or equal to the value of the macro **WEOF**. If the parameter has any other value, the behavior is undefined.

### Parameters

*wc* Specifies the value to be tested.

### Return Values

The **iswblank** subroutine returns a nonzero value if the *wc* parameter is a blank wide-character code; otherwise, it returns a 0.

## Related Information

“iswalnum, iswalph, iswcntrl, iswdigit, iswgraph, iswlower, iswpunct, iswspace, iswupper, or iswxdigit Subroutine” on page 572 and “iswctype or is\_wctype Subroutine.”

setlocale Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

**wctype.h** in *AIX 5L Version 5.3 Files Reference*.

---

## iswctype or is\_wctype Subroutine

### Purpose

Determines properties of a wide character.

### Library

Standard C Library (**libc. a**)

### Syntax

```
#include <wchar.h>
```

```
int iswctype ( WC, Property)
wint_t WC;
wctype_t Property;
```

```
int is_wctype ( WC, Property)
wint_t WC;
wctype_t Property;
```

### Description

The **iswctype** subroutine tests the wide character specified by the *WC* parameter to determine if it has the property specified by the *Property* parameter. The **iswctype** subroutine is defined for the wide-character null value and for values in the character range of the current code set, defined in the current locale. The **is\_wctype** subroutine is identical to the **iswctype** subroutine.

The **iswctype** subroutine adheres to X/Open Portability Guide Issue 5.

### Parameters

*WC* Specifies the wide character to be tested.  
*Property* Specifies the property for which to test.

### Return Values

If the *WC* parameter has the property specified by the *Property* parameter, the **iswctype** subroutine returns a nonzero value. If the value specified by the *WC* parameter does not have the property specified by the *Property* parameter, the **iswctype** subroutine returns a value of zero. If the value specified by the *WC* parameter is not in the subroutine's domain, the result is undefined. If the value specified by the *Property* parameter is not valid (that is, not obtained by a call to the **wctype** subroutine, or the *Property* parameter has been invalidated by a subsequent call to the **setlocale** subroutine that has affected the **LC\_CTYPE** category), the result is undefined.

## Related Information

The “iswalnum, iswalpha, iswcntrl, iswdigit, iswgraph, iswlower, iswprint, iswpunct, iswspace, iswupper, or iswxdigit Subroutine” on page 572.

Subroutines, Example Programs, and Libraries, Wide Character Classification Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## jcode Subroutines

### Purpose

Perform string conversion on 8-bit processing codes.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <jcode.h>
```

```
char *jistosj( String1, String2)
```

```
char *String1, *String2;
```

```
char *jistouj(String1, String2)
```

```
char *String1, *String2;
```

```
char *sjtojis(String1, String2)
```

```
char *String1, *String2;
```

```
char *sjtouj(String1, String2)
```

```
char *String1, *String2;
```

```
char *ujtojis(String1, String2)
```

```
char *String1, *String2;
```

```
char *ujtosj(String1, String2)
```

```
char *String1, *String2;
```

```
char *cjistosj(String1, String2)
```

```
char *String1, *String2;
```

```
char *cjistouj(String1, String2)
```

```
char *String1, *String2;
```

```
char *csjtojis(String1, String2)
```

```
char *String1, *String2;
```

```
char *csjtouj(String1, String2)
```

```
char *String1, *String2;
```

```
char *cujtojis(String1, String2)
```

```
char *String1, *String2;
```

```
char *cujtosj(String1, String2)
```

```
char *String1, *String2;
```

### Description

The **jistosj**, **jistouj**, **sjtojis**, **sjtouj**, **ujtojis**, and **ujtosj** subroutines perform string conversion on 8-bit processing codes. The *String2* parameter is converted and the converted string is stored in the *String1* parameter. The overflow of the *String1* parameter is not checked. Also, the *String2* parameter must be a valid string. Code validation is not permitted.

The **jistosj** subroutine converts JIS to SJIS. The **jistouj** subroutine converts JIS to UJIS. The **sjtojis** subroutine converts SJIS to JIS. The **sjtouj** subroutine converts SJIS to UJIS. The **ujtojis** subroutine converts UJIS to JIS. The **ujtosj** subroutine converts UJIS to SJIS.

The **cjistosj**, **cjistouj**, **csjtojis**, **csjtouj**, **cujtojis**, and **cujtosj** macros perform code conversion on 8-bit processing JIS Kanji characters. A character is removed from the *String2* parameter, and its code is converted and stored in the *String1* parameter. The *String1* parameter is returned. The validity of the *String2* parameter is not checked.

The **cjistosj** macro converts from JIS to SJIS. The **cjistouj** macro converts from JIS to UJIS. The **csjtojis** macro converts from SJIS to JIS. The **csjtouj** macro converts from SJIS to UJIS. The **cujtojis** macro converts from UJIS to JIS. The **cujtosj** macro converts from UJIS to SJIS.

## Parameters

*String1*        Stores converted string or code.  
*String2*        Stores string or code to be converted.

## Related Information

The “Japanese conv Subroutines” and “Japanese ctype Subroutines” on page 578.

List of String Manipulation Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview for Programming in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## Japanese conv Subroutines

### Purpose

Translates predefined Japanese character classes.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <ctype.h>
int atojis ( Character)
int Character;
```

```
int jistoa (Character)
int Character;
```

```
int _atojis (Character)
int Character;
```

```
int _jistoa (Character)
int Character;
```



```
int tojupper (Character)
int Character;
```

```
int tojlower (Character)
int Character;
```

```
int _tojupper (Character)
int Character;
```

```
int _tojlower (Character)
int Character;
```

```
int toujis (Character)
int Character;
```

```
int kutentojis (Character)
int Character;
```

```
int tojhira (Character)
int Character;
```

```
int tojkata (Character)
int Character;
```

## Description

When running the operating system with Japanese Language Support on your system, the legal value of the *Character* parameter is in the range from 0 to **NLCOLMAX**.

The **jistoa** subroutine converts an SJIS ASCII equivalent to the corresponding ASCII equivalent. The **atojis** subroutine converts an ASCII character to the corresponding SJIS equivalent. Other values are returned unchanged.

The **\_jistoa** and **\_atojis** routines are macros that function like the **jistoa** and **atojis** subroutines, but are faster and have no error checking function.

The **tojlower** subroutine converts a SJIS uppercase letter to the corresponding SJIS lowercase letter. The **tojupper** subroutine converts an SJIS lowercase letter to the corresponding SJIS uppercase letter. All other values are returned unchanged.

The **\_tojlower** and **\_tojupper** routines are macros that function like the **tojlower** and **tojupper** subroutines, but are faster and have no error-checking function.

The **toujis** subroutine sets all parameter bits that are not 16-bit SJIS code to 0.

The **kutentojis** subroutine converts a kuten code to the corresponding SJIS code. The **kutentojis** routine returns 0 if the given kuten code is invalid.

The **tojhira** subroutine converts an SJIS katakana character to its SJIS hiragana equivalent. Any value that is not an SJIS katakana character is returned unchanged.

The **tojkata** subroutine converts an SJIS hiragana character to its SJIS katakana equivalent. Any value that is not an SJIS hiragana character is returned unchanged.

The **\_tojhira** and **\_tojkata** subroutines attempt the same conversions without checking for valid input.

For all functions except the **toujis** subroutine, the out-of-range parameter values are returned without conversion.

## Parameters

<i>Character</i>	Character to be converted.
<i>Pointer</i>	Pointer to the escape sequence.
<i>CharacterPointer</i>	Pointer to a single <b>NLchar</b> data type.

## Related Information

The “ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines” on page 208, “conv Subroutines” on page 183, “getc, getchar, fgetc, or getw Subroutine” on page 347, “getwc, fgetwc, or getwchar Subroutine” on page 477, and **setlocale** subroutine.

List of Character Manipulation Subroutines and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*

---

## Japanese ctype Subroutines

### Purpose

Classify characters.

### Library

Standard Character Library (**libc.a**)

### Syntax

```
#include <ctype.h>
```

```
int isjalpha ( Character )  
int Character;
```

```
int isjupper ( Character )  
int Character;
```

```
int isjlower ( Character )  
int Character;
```

```
int isjbytekana ( Character )  
int Character;
```

```
int isjdigit ( Character )  
int Character;
```

```
int isjxdigit ( Character )  
int Character;
```

```
int isjalnum ( Character )  
int Character;
```

```
int isjspace (Character)
```

```
int Character;
```

```
int isjpunct (Character)
```

```
int Character;
```

```
int isjparen (Character)
```

```
int Character;
```

```
int isparent (Character)
```

```
int Character;
```

```
int isjprint (Character)
```

```
int Character;
```

```
int isjgraph (Character)
```

```
int Character;
```

```
int isjis (Character)
```

```
int Character;
```

```
int isjhira (wc)
```

```
wchar_t wc;
```

```
int isjkanji (wc)
```

```
wchar_wc;
```

```
int isjkata (wc)
```

```
wchar_t wc;
```

## Description

The **Japanese ctype** subroutines classify character-coded integer values specified in a table. Each of these subroutines returns a nonzero value for True and 0 for False.

## Parameters

*Character*          Character to be tested.

## Return Values

The **isjprint** and **isjgraph** subroutines return a 0 value for user-defined characters.

## Related Information

The “ctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines” on page 208, and **setlocale** subroutine.

List of Character Manipulation Services and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## kill or killpg Subroutine

### Purpose

Sends a signal to a process or to a group of processes.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <signal.h>
```

```
int kill(
    Process,
    Signal)
pid_t Process;
int Signal;
```

```
killpg(
    ProcessGroup, Signal)
int ProcessGroup, Signal;
```

### Description

The **kill** subroutine sends the signal specified by the *Signal* parameter to the process or group of processes specified by the *Process* parameter.

To send a signal to another process, either the real or the effective user ID of the sending process must match the real or effective user ID of the receiving process, and the calling process must have root user authority.

The processes that have the process IDs of 0 and 1 are special processes and are sometimes referred to here as *proc0* and *proc1*, respectively.

Processes can send signals to themselves.

**Note:** Sending a signal does not imply that the operation is successful. All signal operations must pass the access checks prescribed by each enforced access control policy on the system.

The following interface is provided for BSD Compatibility:

```
killpg(ProcessGroup, Signal)
int ProcessGroup; Signal;
```

This interface is equivalent to:

```
if (ProcessGroup < 0)
{
    errno = ESRCH;
    return (-1);
}
return (kill(-ProcessGroup, Signal));
```

## Parameters

<i>Process</i>	<p>Specifies the ID of a process or group of processes.</p> <p>If the <i>Process</i> parameter is greater than 0, the signal specified by the <i>Signal</i> parameter is sent to the process identified by the <i>Process</i> parameter.</p> <p>If the <i>Process</i> parameter is 0, the signal specified by the <i>Signal</i> parameter is sent to all processes, excluding <i>proc0</i> and <i>proc1</i>, whose process group ID matches the process group ID of the sender.</p> <p>If the value of the <i>Process</i> parameter is a negative value other than -1 and if the calling process passes the access checks for the process to be signaled, the signal specified by the <i>Signal</i> parameter is sent to all the processes, excluding <i>proc0</i> and <i>proc1</i>. If the user ID of the calling process has root user authority, all processes, excluding <i>proc0</i> and <i>proc1</i>, are signaled.</p> <p>If the value of the <i>Process</i> parameter is a negative value other than -1, the signal specified by the <i>Signal</i> parameter is sent to all processes having a process group ID equal to the absolute value of the <i>Process</i> parameter.</p> <p>If the value of the <i>Process</i> parameter is -1, the signal specified by the <i>Signal</i> parameter is sent to all processes which the process has permission to send that signal.</p>
<i>Signal</i>	<p>Specifies the signal. If the <i>Signal</i> parameter is a null value, error checking is performed but no signal is sent. This parameter is used to check the validity of the <i>Process</i> parameter.</p>
<i>ProcessGroup</i>	<p>Specifies the process group.</p>

## Return Values

Upon successful completion, the **kill** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **kill** subroutine is unsuccessful and no signal is sent if one or more of the following are true:

<b>EINVAL</b>	The <i>Signal</i> parameter is not a valid signal number.
<b>EINVAL</b>	The <i>Signal</i> parameter specifies the <b>SIGKILL</b> , <b>SIGSTOP</b> , <b>SIGTSTP</b> , or <b>SIGCONT</b> signal, and the <i>Process</i> parameter is 1 ( <i>proc1</i> ).
<b>ESRCH</b>	No process can be found corresponding to that specified by the <i>Process</i> parameter.
<b>EPERM</b>	The real or effective user ID does not match the real or effective user ID of the receiving process, or else the calling process does not have root user authority.

## Related Information

The **getpid**, **getpgrp**, or **getppid** (“getpid, getpgrp, or getppid Subroutine” on page 407) subroutine, **setpgid** or **setpgrp** subroutine, **sigaction**, **sigvec**, or **signal** subroutine.

The **kill** command.

Signal Management in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* provides more information about signal management in multi-threaded processes.

---

## kleenup Subroutine

### Purpose

Cleans up the run-time environment of a process.

## Library

### Syntax

```
int kleanup( FileDescriptor, SigIgn, SigKeep)
int FileDescriptor;
int SigIgn[ ];
int SigKeep[ ];
```

### Description

The **kleanup** subroutine cleans up the run-time environment for a trusted process by:

- Closing unnecessary file descriptors.
- Resetting the alarm time.
- Resetting signal handlers.
- Clearing the value of the **real directory read** flag described in the **ulimit** subroutine.
- Resetting the **ulimit** value, if it is less than a reasonable value (8192).

### Parameters

<i>FileDescriptor</i>	Specifies a file descriptor. The <b>kleanup</b> subroutine closes all file descriptors greater than or equal to the <i>FileDescriptor</i> parameter.
<i>SigIgn</i>	Points to a list of signal numbers. If these are nonnull values, this list is terminated by 0s. Any signals specified by the <i>SigIgn</i> parameter are set to <b>SIG_IGN</b> . The handling of all signals not specified by either this list or the <i>SigKeep</i> list are set to <b>SIG_DFL</b> . Some signals cannot be reset and are left unchanged.
<i>SigKeep</i>	Points to a list of signal numbers. If these are nonnull values, this list is terminated by 0s. The handling of any signals specified by the <i>SigKeep</i> parameter is left unchanged. The handling of all signals not specified by either this list or the <i>SigIgn</i> list are set to <b>SIG_DFL</b> . Some signals cannot be reset and are left unchanged.

### Return Values

The **kleanup** subroutine is always successful and returns a value of 0. Errors in closing files are not reported. It is not an error to attempt to modify a signal that the process is not allowed to handle.

### Related Information

The **ulimit** subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## knlist Subroutine

### Purpose

Translates names to addresses in the running system.

### Syntax

```
#include <nlist.h>
```

```

int knlist( NList, NumberOfElements, Size)
struct nlist *NList;
int NumberOfElements;
int Size;

```

## Description

The **knlist** subroutine allows a program to examine the list of symbols exported by kernel routines to other kernel modules.

The first field in the **nlist** structure is an input parameter to the **knlist** subroutine. The **n\_value** field is modified by the **knlist** subroutine, and all the others remain unchanged. The **nlist** structure consists of the following fields:

char *n_name	Specifies the name of the symbol whose attributes are to be retrieved.
long n_value	Indicates the virtual address of the object. This will also be the offset when using segment descriptor 0 as the extension parameter of the <b>readx</b> or <b>writex</b> subroutines against the <b>/dev/mem</b> file.

If the name is not found, all fields, other than **n\_name**, are set to 0.

The **nlist.h** file is automatically included by the **a.out.h** file for compatibility. However, do not include the **a.out.h** file if you only need the information necessary to use the **knlist** subroutine. If you do include the **a.out.h** file, follow the **#include** statement with the line:

```
#undef n_name
```

### Notes:

1. If both the **nlist.h** and **netdb.h** files are to be included, the **netdb.h** file should be included before the **nlist.h** file in order to avoid a conflict with the **n\_name** structure member. Likewise, if both the **a.out.h** and **netdb.h** files are to be included, the **netdb.h** file should be included before the **a.out.h** file to avoid a conflict with the **n\_name** structure.
2. If the **netdb.h** file and either the **nlist.h** or **syms.h** file are included, the **n\_name** field will be defined as **\_n.\_n\_name**. This definition allows you to access the **n\_name** field in the **nlist** or **syment** structure. If you need to access the **n\_name** field in the **netent** structure, undefine the **n\_name** field by entering:

```
#undef n_name
```

before accessing the **n\_name** field in the **netent** structure. If you need to access the **n\_name** field in a **syment** or **nlist** structure after undefining it, redefine the **n\_name** field with:

```
#define n_name _n._n_name
```

## Parameters

<i>NList</i>	Points to an array of <b>nlist</b> structures.
<i>NumberOfElements</i>	Specifies the number of structures in the array of <b>nlist</b> structures.
<i>Size</i>	Specifies the size of each structure.

## Return Values

Upon successful completion, **knlist** returns a value of 0. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **knlist** subroutine fails when the following is true:

**EFAULT** Indicates that the *NList* parameter points outside the limit of the array of **nlist** structures.

---

## kpystate Subroutine

### Purpose

Returns the status of a process.

### Syntax

```
kpystate (pid)
pid_t pid;
```

### Description

The **kpystate** subroutine returns the state of a process specified by the *pid* parameter. The **kpystate** subroutine can only be called by a process.

### Parameters

*pid* Specifies the product ID.

### Return Values

If the *pid* parameter is not valid, KP\_NOTFOUND is returned. If the *pid* parameter is valid, the following settings in the process state determine what is returned:

SNONE	Return KP_NOTFOUND.
SIDL	Return KP_INITING.
SZOMB	Return KP_EXITING, also if SEXIT in pv_flag.
SSTOP	Return KP_STOPPED.

Otherwise the *pid* is alive and KP\_ALIVE is returned.

### Error Codes

---

## \_lazySetErrorHandler Subroutine

### Purpose

Installs an error handler into the lazy loading runtime system for the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/ldr.h>
#include <sys/errno.h>
typedef void (*_handler_t(
char *_module,
char *_symbol,
unsigned int _errVal ))();
handler_t *_lazySetErrorHandler(err_handler)
handler_t *err_handler;
```



## Description

This function allows a process to install a custom error handler to be called when a lazy loading reference fails to find the required module or function. This function should only be used when the main program or one of its dependent modules was linked with the **-blazy** option. To call **\_lazySetErrorHandler** from a module that is not linked with the **-blazy** option, you must use the **-lrtl** option. If you use **-blazy**, you do not need to specify **-lrtl**.

This function is not thread safe. The calling program should ensure that **\_lazySetErrorHandler** is not called by multiple threads at the same time.

The user-supplied error handler may print its own error message, provide a substitute function to be used in place of the called function, or call the **longjmp** subroutine. To provide a substitute function that will be called instead of the originally referenced function, the error handler should return a pointer to the substitute function. This substitute function will be called by all subsequent calls to the intended function from the same module. If the value returned by the error handler appears to be invalid (for example, a NULL pointer), the default error handler will be used.

Each calling module resolves its lazy references independent of other modules. That is, if module A and B both call **foo** subroutine in module C, but module C does not export **foo** subroutine, the error handler will be called once when **foo** subroutine is called for the first time from A, and once when **foo** subroutine is called for the first time from B.

The default lazy loading error handler will print a message containing: the name of module that the program required; the name of the symbol being accessed; and the error value generated by the failure. Since the default handler considers a lazy load error to be fatal, the process will exit with a status of 1.

During execution of a program that utilizes lazy loading, there are a few conditions that may cause an error to occur. In all cases the current error handler will be called.

1. The referenced module (which is to be loaded upon function invocation) is unavailable or cannot be loaded. The *errVal* parameter will probably indicate the reason for the error if a system call failed.
2. A function is referenced, but the loaded module does not contain a definition for the function. In this case, *errVal* parameter will be **EINVAL**.

Some possibilities as to why either of these errors might occur:

1. The **LIBPATH** environment variable may contain a set of search paths that cause the application to load the wrong version of a module.
2. A module has been changed and no longer provides the same set of symbols that it did when the application was built.
3. The **load** subroutine fails due to a lack of resources available to the process.

## Parameters

*err\_handler* A pointer to the new error handler function. The new function should accept 3 arguments:

- module* The name of the referenced module.
- symbol* The name of the function being called at the time the failure occurred.
- errVal* The value of **errno** at the time the failure occurred, if a system call used to load the module fails. For other failures, *errval* may be **EINVAL** or **ENOMEM**.

Note that the value of *module* or *symbol* may be NULL if the calling module has somehow been corrupted.

If the *err\_handler* parameter is NULL, the default error handler is restored.

## Return Value

The function returns a pointer to the previous user-supplied error handler, or NULL if the default error handler was in effect.

## Related Information

The **load** (“load and loadAndInit Subroutines” on page 727) subroutine.

The **ld** command.

The Shared Library Overview and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts*.

The Shared Library and Lazy Loading in *AIX 5L Version 5.3 General Programming Concepts*.

---

## l3tol or ltol3 Subroutine

### Purpose

Converts between 3-byte integers and long integers.

### Library

Standard C Library (**libc.a**)

### Syntax

```
void l3tol ( LongPointer, CharacterPointer, Number)
```

```
long *LongPointer;
```

```
char *CharacterPointer;
```

```
int Number;
```

```
void ltol3 (CharacterPointer, LongPointer, Number)
```

```
char *CharacterPointer;
```

```
long *LongPointer;
```

```
int Number;
```

### Description

The **l3tol** subroutine converts a list of the number of 3-byte integers specified by the *Number* parameter packed into a character string pointed to by the *CharacterPointer* parameter into a list of long integers pointed to by the *LongPointer* parameter.

The **ltol3** subroutine performs the reverse conversion, from long integers (the *LongPointer* parameter) to 3-byte integers (the *CharacterPointer* parameter).

These functions are useful for file system maintenance where the block numbers are 3 bytes long.

### Parameters

*LongPointer*

Specifies the address of a list of long integers.

*CharacterPointer*

Specifies the address of a list of 3-byte integers.

*Number*

Specifies the number of list elements to convert.

## Related Information

The **filsys.h** file format.

## I64a\_r Subroutine

### Purpose

Converts base-64 long integers to strings.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
#include <stdlib.h>
```

```
int I64a_r (Convert, Buffer, Length)
long Convert;
char * Buffer;
int Length;
```

### Description

The **I64a\_r** subroutine converts a given long integer into a base-64 string.

Programs using this subroutine must link to the **libpthread.a** library.

For base-64 characters, the following ASCII characters are used:

Character	Description
.	Represents 0.
/	Represents 1.
0 -9	Represents the numbers 2-11.
A-Z	Represents the numbers 12-37.
a-z	Represents the numbers 38-63.

The **I64a\_r** subroutine places the converted base-64 string in the buffer pointed to by the *Buffer* parameter.

### Parameters

<i>Convert</i>	Specifies the long integer that is to be converted into a base-64 ASCII string.
<i>Buffer</i>	Specifies a working buffer to hold the converted long integer.
<i>Length</i>	Specifies the length of the <i>Buffer</i> parameter.

### Return Values

0	Indicates that the subroutine was successful.
-1	Indicates that the subroutine was not successful. If the <b>I64a_r</b> subroutine is not successful, the <b>errno</b> global variable is set to indicate the error.

## Error Codes

If the **I64a\_r** subroutine is not successful, it returns the following error code:

**EINVAL**                    The *Buffer* parameter value is invalid or too small to hold the resulting ASCII string.

## Related Information

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

List of Multithread Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## LAPI\_Addr\_get Subroutine

### Purpose

Retrieves a function address that was previously registered using **LAPI\_Addr\_set**.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Addr_get(hndl, addr, addr_hndl)
lapi_handle_t hndl;
void          **addr;
int           addr_hndl;
```

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_ADDR_GET(hndl, addr, addr_hndl, ierror)
INTEGER hndl
INTEGER (KIND=LAPI_ADDR_TYPE) :: addr
INTEGER addr_hndl
INTEGER ierror
```

### Description

**Type of call:** local address manipulation

Use this subroutine to get the pointer that was previously registered with LAPI and is associated with the index *addr\_hndl*. The value of *addr\_hndl* must be in the range **1** <= *addr\_hndl* < **LOC\_ADDRTBL\_SZ**.

### Parameters

#### INPUT

*hndl*                    Specifies the LAPI handle.

*addr\_hndl*              Specifies the index of the function address to retrieve. You should have previously registered the address at this index using **LAPI\_Addr\_set**. The value of this parameter must be in the range **1** <= *addr\_hndl* < **LOC\_ADDRTBL\_SZ**.

#### OUTPUT

*addr* Returns a function address that the user registered with LAPI.  
*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To retrieve a header handler address that was previously registered using **LAPI\_Addr\_set**:

```
lapi_handle_t hndl; /* the LAPI handle */
void **addr; /* the address to retrieve */
int addr_hndl; /* the index returned from LAPI_Addr_set */

:

addr_hndl = 1;
LAPI_Addr_get(hndl, &addr, addr_hndl);

/* addr now contains the address that was previously registered */
/* using LAPI_Addr_set */
```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

### **LAPI\_ERR\_ADDR\_HNDL\_RANGE**

Indicates that the value of *addr\_hndl* is not in the range  $1 \leq \text{addr\_hndl} < \text{LOC\_ADDRTBL\_SZ}$ .

### **LAPI\_ERR\_HNDL\_INVALID**

Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

### **LAPI\_ERR\_RET\_PTR\_NULL**

Indicates that the value of the *addr* pointer is NULL (in C) or that the value of *addr* is **LAPI\_ADDR\_NULL** (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Addr\_set**, **LAPI\_Qenv**

---

## LAPI\_Addr\_set Subroutine

### Purpose

Registers the address of a function.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Addr_set(hndl, addr, addr_hndl)
lapi_handle_t hndl;
void **addr;
int addr_hndl;
```

## FORTRAN Syntax

```
include 'lapif.h'

LAPI_ADDR_SET(hndl, addr, addr_hndl, ierror)
INTEGER hndl
INTEGER (KIND=LAPI_ADDR_TYPE) :: addr
INTEGER addr_hndl
INTEGER ierror
```

## Description

**Type of call:** local address manipulation

Use this subroutine to register the address of a function (*addr*). LAPI maintains the function address in an internal table. The function address is indexed at location *addr\_hndl*. In subsequent LAPI calls, *addr\_hndl* can be used in place of *addr*. The value of *addr\_hndl* must be in the range  $1 \leq \text{addr\_hndl} < \text{LOC\_ADDRTBL\_SZ}$ .

For active message communication, you can use *addr\_hndl* in place of the corresponding header handler address. LAPI only supports this indexed substitution for remote header handler addresses (but not other remote addresses, such as target counters or base data addresses). For these other types of addresses, the actual address value must be passed to the API call.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*addr* Specifies the address of the function handler that the user wants to register with LAPI.

*addr\_hndl* Specifies a user function address that can be passed to LAPI calls in place of a header handler address. The value of this parameter must be in the range  $1 \leq \text{addr\_hndl} < \text{LOC\_ADDRTBL\_SZ}$ .

### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To register a header handler address:

```
lapi_handle_t hndl; /* the LAPI handle */
void *addr; /* the remote header handler address */
int addr_hndl; /* the index to associate */

:

addr = my_func;
addr_hndl = 1;
LAPI_Addr_set(hndl, addr, addr_hndl);

/* addr_hndl can now be used in place of addr in LAPI_Amsend, */
/* LAPI_Amsendv, and LAPI_Xfer calls */

:
```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

## LAPI\_ERR\_ADDR\_HNDL\_RANGE

Indicates that the value of *addr\_hndl* is not in the range  $1 \leq \text{addr\_hndl} < \text{LOC\_ADDRTBL\_SZ}$ .

## LAPI\_ERR\_HNDL\_INVALID

Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Addr\_get**, **LAPI\_Amsend**, **LAPI\_Amsendv**, **LAPI\_Qenv**, **LAPI\_Xfer**

---

## LAPI\_Address Subroutine

### Purpose

Returns an unsigned long value for a specified user address.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Address(my_addr, ret_addr)
void *my_addr;
ulong *ret_addr;
```

**Note:** This subroutine is meant to be used by FORTRAN programs. The C version of **LAPI\_Address** is provided for compatibility purposes only.

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_ADDRESS(my_addr, ret_addr, ierror)
INTEGER (KIND=any_type) :: my_addr
INTEGER (KIND=LAPI_ADDR_TYPE) :: ret_addr
INTEGER ierror
```

where:

*any\_type*      Is any FORTRAN datatype. This type declaration has the same meaning as the type **void \*** in C.

### Description

**Type of call:** local address manipulation

Use this subroutine in FORTRAN programs when you need to store specified addresses in an array. In FORTRAN, the concept of address (&) does not exist as it does in C. **LAPI\_Address** provides FORTRAN programmers with this function.

### Parameters

#### INPUT

*my\_addr* Specifies the address to convert. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

## OUTPUT

*ret\_addr* Returns the address that is stored in *my\_addr* as an unsigned long for use in LAPI calls. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## FORTRAN Examples

To retrieve the address of a variable:

```
! Contains the address of the target counter
integer (KIND=LAPI_ADDR_TYPE) :: cntr_addr

! Target Counter
type (LAPI_CNTR_T) :: tgt_cntr

! Return code
integer :: ierror

call LAPI_ADDRESS(tgt_cntr, cntr_addr, ierror)

! cntr_addr now contains the address of tgt_cntr
```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_ORG\_ADDR\_NULL** Indicates that the value of *my\_addr* is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_TGT\_ADDR\_NULL** Indicates that the value of *ret\_addr* is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address\_init**, **LAPI\_Address\_init64**

---

## LAPI\_Address\_init Subroutine

### Purpose

Creates a remote address table.

### Library

Availability Library (**liblapi\_r.a**)



## C Syntax

```
#include <lapi.h>

int LAPI_Address_init(hndl, my_addr, add_tab)
lapi_handle_t hndl;
void          *my_addr;
void          *add_tab[ ];
```

## FORTRAN Syntax

```
include 'lapif.h'

LAPI_ADDRESS_INIT(hndl, my_addr, add_tab, ierror)
INTEGER hndl
INTEGER (KIND=LAPI_ADDR_TYPE) :: my_addr
INTEGER (KIND=LAPI_ADDR_TYPE) :: add_tab(*)
INTEGER ierror
```

## Description

**Type of call:** collective communication (blocking)

**LAPI\_Address\_init** exchanges virtual addresses among tasks of a parallel application. Use this subroutine to create tables of such items as header handlers, target counters, and data buffer addresses.

**LAPI\_Address\_init** is a *collective call* over the LAPI handle *hndl*, which fills the table *add\_tab* with the virtual address entries that each task supplies. Collective calls must be made in the same order at all participating tasks.

The addresses that are stored in the table *add\_tab* are passed in using the *my\_addr* parameter. Upon completion of this call, *add\_tab[i]* contains the virtual address entry that was provided by task *i*. The array is opaque to the user.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*my\_addr* Specifies the entry supplied by each task. The value of this parameter can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

### OUTPUT

*add\_tab* Specifies the address table containing the addresses that are to be supplied by all tasks. *add\_tab* is an array of pointers, the size of which is greater than or equal to **NUM\_TASKS**. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To collectively transfer target counter addresses for use in a communication API call, in which all nodes are either 32-bit or 64-bit:

```
lapi_handle_t hndl;          /* the LAPI handle          */
void          *addr_tbl[NUM_TASKS]; /* the table for all tasks' addresses */
lapi_cntr_t   tgt_cntr;      /* the target counter      */
:
:
LAPI_Address_init(hndl, (void *)&tgt_cntr, addr_tbl);
```

```

/* for communication with task t, use addr_tbl[t] */
/* as the address of the target counter          */
:

```

For a combination of 32-bit and 64-bit nodes, use **LAPI\_Address\_init64**.

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_COLLECTIVE_PSS</b>	Indicates that a collective call was made while in persistent subsystem (PSS) mode.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hdl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_RET_PTR_NULL</b>	Indicates that the value of the <i>add_tab</i> pointer is NULL (in C) or that the value of <i>add_tab</i> is <b>LAPI_ADDR_NULL</b> (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address**, **LAPI\_Address\_init64**

---

## LAPI\_Address\_init64 Subroutine

### Purpose

Creates a 64-bit remote address table.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```

#include <lapi.h>

int LAPI_Address_init64(hndl, my_addr, add_tab)
lapi_handle_t hndl;
lapi_long_t my_addr;
lapi_long_t *add_tab;

```

### FORTRAN Syntax

```

include 'lapif.h'

LAPI_ADDRESS_INIT64(hndl, my_addr, add_tab, ierror)
INTEGER hndl
INTEGER (KIND=LAPI_ADDR_TYPE) :: my_addr
INTEGER (KIND=LAPI_LONG_LONG_TYPE) :: add_tab(*)
INTEGER ierror

```

### Description

**Type of call:** collective communication (blocking)

**LAPI\_Address\_init64** exchanges virtual addresses among a mixture of 32-bit and 64-bit tasks of a parallel application. Use this subroutine to create 64-bit tables of such items as header handlers, target counters, and data buffer addresses.

**LAPI\_Address\_init64** is a *collective call* over the LAPI handle *hndl*, which fills the 64-bit table *add\_tab* with the virtual address entries that each task supplies. Collective calls must be made in the same order at all participating tasks.

The addresses that are stored in the table *add\_tab* are passed in using the *my\_addr* parameter. Upon completion of this call, *add\_tab[i]* contains the virtual address entry that was provided by task *i*. The array is opaque to the user.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*my\_addr* Specifies the address entry that is supplied by each task. The value of this parameter can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). To ensure 32-bit/64-bit interoperability, it is passed as a **lapi\_long\_t** type in C.

### OUTPUT

*add\_tab* Specifies the 64-bit address table that contains the 64-bit values supplied by all tasks. *add\_tab* is an array of type **lapi\_long\_t** (in C) or **LAPI\_LONG\_LONG\_TYPE** (in FORTRAN). The size of *add\_tab* is greater than or equal to **NUM\_TASKS**. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To collectively transfer target counter addresses for use in a communication API call with a mixed task environment (any combination of 32-bit and 64-bit):

```
lapi_handle_t hndl;           /* the LAPI handle           */
lapi_long_t  addr_tbl[NUM_TASKS]; /* the table for all tasks' addresses */
lapi_long_t  tgt_cntr;       /* the target counter       */
:
LAPI_Address_init64(hndl, (lapi_long_t)&tgt_cntr, addr_tbl);

/* For communication with task t, use addr_tbl[t] as the address */
/* of the target counter. For mixed (32-bit and 64-bit) jobs,    */
/* use the LAPI_Xfer subroutine for communication.              */
```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_COLLECTIVE\_PSS** Indicates that a collective call was made while in persistent subsystem (PSS) mode.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_RET\_PTR\_NULL** Indicates that the value of the *add\_tab* pointer is NULL (in C) or that the value of *add\_tab* is **LAPI\_ADDR\_NULL** (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address**, **LAPI\_Address\_init**, **LAPI\_Xfer**

---

## LAPI\_Amsend Subroutine

### Purpose

Transfers a user message to a remote task, obtaining the target address on the remote task from a user-specified header handler.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
typedef void (compl_hdlr_t) (hdl, user_info);
```

```
lapi_handle_t *hdl;      /* pointer to LAPI context passed in from LAPI_Amsend */  
void          *user_info; /* buffer (user_info) pointer passed in          */  
                /* from header handler (void *(hdr_hdlr_t))      */
```

```
typedef void *(hdr_hdlr_t)(hdl, uhdr, uhdr_len, msg_len, comp_h, user_info);
```

```
lapi_handle_t *hdl;      /* pointer to LAPI context passed in from LAPI_Amsend */  
void          *uhdr;     /* uhdr passed in from LAPI_Amsend          */  
uint          *uhdr_len; /* uhdr_len passed in from LAPI_Amsend     */  
ulong         *msg_len;  /* udata_len passed in from LAPI_Amsend    */  
compl_hdlr_t **comp_h;  /* function address of completion handler  */  
                /* (void (compl_hdlr_t)) that needs to be filled */  
                /* out by this header handler function.      */  
void          **user_info; /* pointer to the parameter to be passed  */  
                /* in to the completion handler                */
```

```
int LAPI_Amsend(hdl, tgt, hdr_hdl, uhdr, uhdr_len, udata, udata_len,  
                tgt_cntr, org_cntr, compl_cntr)
```

```
lapi_handle_t hdl;  
uint          tgt;  
void          *hdr_hdl;  
void          *uhdr;  
uint          uhdr_len;  
void          *udata;  
ulong         udata_len;  
lapi_cntr_t  *tgt_cntr;  
lapi_cntr_t  *org_cntr;  
lapi_cntr_t  *compl_cntr;
```

### FORTTRAN Syntax

```
include 'lapif.h'
```

```
INTEGER SUBROUTINE COMPL_H (hdl, user_info)  
INTEGER hdl  
INTEGER user_info
```

```

INTEGER FUNCTION HDR_HDL (hdl, uhdr, uhdr_len, msg_len, comp_h, user_info)
INTEGER hdl
INTEGER uhdr
INTEGER uhdr_len
INTEGER (KIND=LAPI_LONG_TYPE) :: msg_len
EXTERNAL INTEGER FUNCTION comp_h
TYPE (LAPI_ADDR_T) :: user_info

LAPI_AMSEND(hdl, tgt, hdr_hdl, uhdr, uhdr_len, udata, udata_len,
           tgt_cntr, org_cntr, cpl_cntr, ierror)
INTEGER hdl
INTEGER tgt
EXTERNAL INTEGER FUNCTION hdr_hdl
INTEGER uhdr
INTEGER uhdr_len
TYPE (LAPI_ADDR_T) :: udata
INTEGER (KIND=LAPI_LONG_TYPE) :: udata_len
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_cntr
TYPE (LAPI_CNTR_T) :: org_cntr
TYPE (LAPI_CNTR_T) :: cpl_cntr
INTEGER ierror

```

## Description

**Type of call:** point-to-point communication (non-blocking)

Use this subroutine to transfer data to a target task, where it is desirable to run a handler on the target task before message delivery begins or after delivery completes. **LAPI\_Amsend** allows the user to provide a header handler and optional completion handler. The header handler is used to specify the target buffer address for writing the data, eliminating the need to know the address on the origin task when the subroutine is called.

User data (*uhdr* and *udata*) are sent to the target task. Once these buffers are no longer needed on the origin task, the origin counter is incremented, which indicates the availability of origin buffers for modification. Using the **LAPI\_Xfer** call with the **LAPI\_AM\_XFER** type provides the same type of transfer, with the option of using a send completion handler instead of the origin counter to specify buffer availability.

Upon arrival of the first data packet at the target, the user's header handler is invoked. Note that a header handler must be supplied by the user because it returns the base address of the buffer in which LAPI will write the data sent from the origin task (*udata*). See *RSCT for AIX 5L: LAPI Programming Guide* for an optimization exception to this requirement that a buffer address be supplied to LAPI for single-packet messages.

The header handler also provides additional information to LAPI about the message delivery, such as the completion handler. **LAPI\_Amsend** and similar calls (such as **LAPI\_Amsendv** and corresponding **LAPI\_Xfer** transfers) also allow the user to specify their own message header information, which is available to the header handler. The user may also specify a completion handler parameter from within the header handler. LAPI will pass the information to the completion handler at execution.

Note that the header handler is run inline by the thread running the LAPI dispatcher. For this reason, the header handler must be non-blocking because no other progress on messages will be made until it returns. It is also suggested that execution of the header handler be simple and quick. The completion handler, on the other hand, is normally enqueued for execution by a separate thread. It is possible to request that the completion handler be run inline. See *RSCT for AIX 5L: LAPI Programming Guide* for more information on inline completion handlers.

If a completion handler was not specified (that is, set to **LAPI\_ADDR\_NULL** in FORTRAN or its pointer set to NULL in C), the arrival of the final packet causes LAPI to increment the target counter on the

remote task and send an internal message back to the origin task. The message causes the completion counter (if it is not NULL in C or **LAPI\_ADDR\_NULL** in FORTRAN) to increment on the origin task.

If a completion handler was specified, the above steps take place after the completion handler returns. To guarantee that the completion handler has executed on the target, you must wait on the completion counter. See *RSCT for AIX 5L: LAPI Programming Guide* for a time-sequence diagram of events in a **LAPI\_Amsend** call.

### User details

As mentioned above, the user must supply the address of a header handler to be executed on the target upon arrival of the first data packet. The signature of the header handler is as follows:

```
void *hdr_hdlr(lapi_handle_t *hdl, void *uhdr, uint *uhdr_len, ulong *msg_len,
               compl_hdlr_t **compl_hdlr, void **user_info);
```

The value returned by the header handler is interpreted by LAPI as an address for writing the user data (*udata*) that was passed to the **LAPI\_Amsend** call. The *uhdr* and *uhdr\_len* parameters are passed by LAPI into the header handler and contain the information passed by the user to the corresponding parameters of the **LAPI\_Amsend** call.

### Use of LAPI\_Addr\_set

Remote addresses are commonly exchanged by issuing a collective **LAPI\_Address\_init** call within a few steps of initializing LAPI. LAPI also provides the **LAPI\_Addr\_set** mechanism, whereby users can register one or more header handler addresses in a table, associating an index value with each address. This index can then be passed to **LAPI\_Amsend** instead of an actual address. On the target side, LAPI will use the index to get the header handler address. Note that, if all tasks use the same index for their header handler, the initial collective communication can be avoided. Each task simply registers its own header handler address using the well-known index. Then, on any **LAPI\_Amsend** calls, the reserved index can be passed to the header handler address parameter.

### Role of the header handler

The user optionally returns the address of a completion handler function through the *compl\_hdlr* parameter and a completion handler parameter through the *user\_info* parameter. The address passed through the *user\_info* parameter can refer to memory containing a datatype defined by the user and then cast to the appropriate type from within the completion handler if desired.

The signature for a user completion handler is as follows:

```
typedef void (compl_hdlr_t)(lapi_handle_t *hdl, void *completion_param);
```

The argument returned by reference through the *user\_info* member of the user's header handler will be passed to the *completion\_param* argument of the user's completion handler. See the **C Examples** for an example of setting the completion handler and parameter in the header handler.

As mentioned above, the value returned by the header handler must be an address for writing the user data sent from the origin task. There is one exception to this rule. In the case of a single-packet message, LAPI passes the address of the packet in the receive FIFO, allowing the entire message to be consumed within the header handler. In this case, the header handler should return NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN) so that LAPI does not copy the message to a target buffer. See *RSCT for AIX 5L: LAPI Programming Guide* for more information (including a sample header handler that uses this method for fast retrieval of a single-packet message).

### Passing additional information through lapi\_return\_info\_t

LAPI allows additional information to be passed to and returned from the header handler by passing a pointer to `lapi_return_info_t` through the `msg_len` argument. On return from a header handler that is invoked by a call to `LAPI_Amsend`, the `ret_flags` member of `lapi_return_info_t` can contain one of these values: `LAPI_NORMAL` (the default), `LAPI_SEND_REPLY` (to run the completion handler inline), or `LAPI_LOCAL_STATE` (no reply is sent). The `dgsp_handle` member of `lapi_return_info_t` should not be used in conjunction with `LAPI_Amsend`.

For a complete description of the `lapi_return_info_t` type, see *RSCT for AIX 5L: LAPI Programming Guide*

### Inline execution of completion handlers

Under normal operation, LAPI uses a separate thread for executing user completion handlers. After the final packet arrives, completion handler pointers are placed in a queue to be handled by this thread. For performance reasons, the user may request that a given completion handler be run inline instead of being placed on this queue behind other completion handlers. This mechanism gives users a greater degree of control in prioritizing completion handler execution for performance-critical messages.

LAPI places no restrictions on completion handlers that are run "normally" (that is, by the completion handler thread). Inline completion handlers should be short and should not block, because no progress can be made while the main thread is executing the handler. The user must use caution with inline completion handlers so that LAPI's internal queues do not fill up while waiting for the handler to complete. I/O operations must not be performed with an inline completion handler.

## Parameters

### INPUT

<i>hndl</i>	Specifies the LAPI handle.
<i>tgt</i>	Specifies the task ID of the target task. The value of this parameter must be in the range $0 \leq \text{tgt} < \text{NUM\_TASKS}$ .
<i>hdr_hdl</i>	Specifies the pointer to the remote header handler function to be invoked at the target. The value of this parameter can take an address handle that has already been registered using <code>LAPI_Addr_set</code> . The value of this parameter cannot be NULL (in C) or <code>LAPI_ADDR_NULL</code> (in FORTRAN).
<i>uhdr</i>	Specifies the pointer to the user header data. This data will be passed to the user header handler on the target. If <i>uhdr_len</i> is 0, The value of this parameter can be NULL (in C) or <code>LAPI_ADDR_NULL</code> (in FORTRAN).
<i>uhdr_len</i>	Specifies the length of the user's header. The value of this parameter must be a multiple of the processor's word size in the range $0 \leq \text{uhdr\_len} \leq \text{MAX\_UHDR\_SZ}$ .
<i>udata</i>	Specifies the pointer to the user data. If <i>udata_len</i> is 0, The value of this parameter can be NULL (in C) or <code>LAPI_ADDR_NULL</code> (in FORTRAN).
<i>udata_len</i>	Specifies the length of the user data in bytes. The value of this parameter must be in the range $0 \leq \text{udata\_len} \leq$ the value of LAPI constant <code>LAPI_MAX_MSG_SZ</code> .

### INPUT/OUTPUT

<i>tgt_cntr</i>	Specifies the target counter address. The target counter is incremented after the completion handler (if specified) completes or after the completion of data transfer. If the value of this parameter is NULL (in C) or <code>LAPI_ADDR_NULL</code> (in FORTRAN), the target counter is not updated.
<i>org_cntr</i>	Specifies the origin counter address (in C) or the origin counter (in FORTRAN). The origin counter is incremented after data is copied out of the origin address (in C) or the origin (in

FORTRAN). If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the origin counter is not updated.

*compl\_cntr* Specifies the counter at the origin that signifies completion of the completion handler. It is updated once the completion handler completes. If no completion handler is specified, the counter is incremented at the completion of message delivery. If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the completion counter is not updated.

## OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_DATA\_LEN** Indicates that the value of *udata\_len* is greater than the value of LAPI constant **LAPI\_MAX\_MSG\_SZ**.

**LAPI\_ERR\_HDR\_HNDLR\_NULL** Indicates that the value of the *hdr\_hdl* passed in is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hdl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_ORG\_ADDR\_NULL** Indicates that the value of the *udata* parameter passed in is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), but the value of *udata\_len* is greater than **0**.

**LAPI\_ERR\_TGT** Indicates that the *tgt* passed in is outside the range of tasks defined in the job.

**LAPI\_ERR\_TGT\_PURGED** Indicates that the subroutine returned early because **LAPI\_Purge\_totask()** was called.

**LAPI\_ERR\_UHDR\_LEN** Indicates that the *uhdr\_len* value passed in is greater than **MAX\_UHDR\_SZ** or is not a multiple of the processor's doubleword size.

**LAPI\_ERR\_UHDR\_NULL** Indicates that the *uhdr* passed in is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), but *uhdr\_len* is not **0**.

## C Examples

To send an active message and then wait on the completion counter:

```
/* header handler routine to execute on target task */
void *hdr_hdlr(lapi_handle_t *hdl, void *uhdr, uint *uhdr_len,
              ulong *msg_len, compl_hdlr_t **compl_hdlr,
              void **user_info)
{
    /* set completion handler pointer and other information */
    /* return base address for LAPI to begin its data copy */
}

{
    lapi_handle_t hdl;          /* the LAPI handle          */
    int          task_id;      /* the LAPI task ID        */
    int          num_tasks;    /* the total number of tasks */
    void        *hdr_hdlr_list[NUM_TASKS]; /* the table of remote header handlers */
    int          buddy;       /* the communication partner */
    lapi_cntr_t  compl_cntr;  /* the completion counter  */
    int          data_buffer[DATA_LEN]; /* the data to transfer    */
}
```



```

.
.
.
/* retrieve header handler addresses */
LAPI_Address_init(hndl, (void *)&hdr_hdlr, hdr_hdlr_list);

/*
** up to this point, all instructions have executed on all
** tasks. we now begin differentiating tasks.
*/
if ( sender ) {          /* origin task */

    /* initialize data buffer, cmpl_cntr, etc. */
    .
    .
    .
    /* synchronize before starting data transfer */
    LAPI_Gfence(hndl);

    LAPI_Amsend(hndl, buddy, (void *)hdr_hdlr_list[buddy], NULL,
                0,&(data_buffer[0]),DATA_LEN*(sizeof(int)),
                NULL, NULL, cmpl_cntr);

    /* Wait on completion counter before continuing. Completion */
    /* counter will update when message completes at target. */

} else {                /* receiver */
    .
    .
    .
    /* to match the origin's synchronization before data transfer */
    LAPI_Gfence(hndl);
}

.
.
.
}

```

For a complete program listing, see *RSCT for AIX 5L: LAPI Programming Guide*. Sample code illustrating the **LAPI\_Amsend** call can be found in the LAPI sample files. See *RSCT for AIX 5L: LAPI Programming Guide* for more information about the sample programs that are shipped with LAPI.

## Location

`/usr/lib/liblapi_r.a`

## Related Information

Subroutines: **LAPI\_Addr\_get**, **LAPI\_Addr\_set**, **LAPI\_Getcntr**, **LAPI\_Msgpoll**, **LAPI\_Qenv**, **LAPI\_Setcntr**, **LAPI\_Waitcntr**, **LAPI\_Xfer**

---

## LAPI\_Amsendv Subroutine

### Purpose

Transfers a user vector to a remote task, obtaining the target address on the remote task from a user-specified header handler.

## Library

Availability Library (**liblapi\_r.a**)

## C Syntax

```
#include <lapi.h>
```

```
typedef void (compl_hdlr_t) (hdl, user_info);
lapi_handle_t *hdl; /* the LAPI handle passed in from LAPI_Amsendv */
void *user_info; /* the buffer (user_info) pointer passed in */
/* from vhdr_hdlr (void *(vhdr_hdlr_t)) */

typedef lapi_vec_t *(vhdr_hdlr_t) (hdl, uhdr, uhdr_len, len_vec, comp_h, uinfo);

lapi_handle_t *hdl; /* pointer to the LAPI handle passed in from LAPI_Amsendv */
void *uhdr; /* uhdr passed in from LAPI_Amsendv */
uint *uhdr_len; /* uhdr_len passed in from LAPI_Amsendv */
ulong *len_vec[ ]; /* vector of lengths passed in LAPI_Amsendv */
compl_hdlr_t **comp_h; /* function address of completion handler */
/* (void (compl_hdlr_t)) that needs to be */
/* filled out by this header handler function */
void **user_info; /* pointer to the parameter to be passed */
/* in to the completion handler */

int LAPI_Amsendv(hdl, tgt, hdr_hdl, uhdr, uhdr_len, org_vec,
                tgt_cntr, org_cntr, compl_cntr);

lapi_handle_t hdl;
uint tgt;
void *hdr_hdl;
void *uhdr;
uint uhdr_len;
lapi_vec_t *org_vec;
lapi_cntr_t *tgt_cntr;
lapi_cntr_t *org_cntr;
lapi_cntr_t *compl_cntr;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
INTEGER SUBROUTINE COMPL_H (hdl, user_info)
```

```
INTEGER hdl
```

```
INTEGER user_info(*)
```

```
INTEGER FUNCTION VHDR_HDL (hdl, uhdr, uhdr_len, len_vec, comp_h, user_info)
```

```
INTEGER hdl
```

```
INTEGER uhdr
```

```
INTEGER uhdr_len
```

```
INTEGER (KIND=LAPI_LONG_TYPE) :: len_vec
```

```
EXTERNAL INTEGER FUNCTION comp_h
```

```
TYPE (LAPI_ADDR_T) :: user_info
```

```
LAPI_AMSENDV(hdl, tgt, hdr_hdl, uhdr, uhdr_len, org_vec,
             tgt_cntr, org_cntr, compl_cntr, ierror)
```

```
INTEGER hdl
```

```
INTEGER tgt
```

```
EXTERNAL INTEGER FUNCTION hdr_hdl
```

```
INTEGER uhdr
```

```
INTEGER uhdr_len
```

```
TYPE (LAPI_VEC_T) :: org_vec
```

```
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_cntr
```

```
TYPE (LAPI_CNTR_T) :: org_cntr
```

```
TYPE (LAPI_CNTR_T) :: compl_cntr
```

```
INTEGER ierror
```

## Description

**Type of call:** point-to-point communication (non-blocking)

**LAPI\_Amsendv** is the vector-based version of the **LAPI\_Amsend** call. You can use it to specify multi-dimensional and non-contiguous descriptions of the data to transfer. Whereas regular LAPI calls allow the specification of a single data buffer address and length, the vector versions allow the specification of a vector of address and length combinations. Additional information is allowed in the data description on the origin task and the target task.

Use this subroutine to transfer a vector of data to a target task, when you want a handler to run on the target task before message delivery begins or after message delivery completes.

To use **LAPI\_Amsendv**, you must provide a header handler, which returns the address of the target vector description that LAPI uses to write the data that is described by the origin vector. The header handler is used to specify the address of the vector description for writing the data, which eliminates the need to know the description on the origin task when the subroutine is called. The header handler is called upon arrival of the first data packet at the target.

Optionally, you can also provide a completion handler. The header handler provides additional information to LAPI about the message delivery, such as the completion handler. You can also specify a completion handler parameter from within the header handler. LAPI passes the information to the completion handler at execution.

With the exception of the address that is returned by the completion handler, the use of counters, header handlers, and completion handlers in **LAPI\_Amsendv** is identical to that of **LAPI\_Amsend**. In both cases, the user header handler returns information that LAPI uses for writing at the target. See **LAPI\_Amsend** for more information. This section presents information that is specific to the vector version of the call (**LAPI\_Amsendv**).

LAPI vectors are structures of type **lapi\_vec\_t**, defined as follows:

```
typedef struct {
    lapi_vectype_t  vec_type;
    uint           num_vecs;
    void           **info;
    ulong          *len;
} lapi_vec_t;
```

*vec\_type* is an enumeration that describes the type of vector transfer, which can be: **LAPI\_GEN\_GENERIC**, **LAPI\_GEN\_IOVECTOR**, or **LAPI\_GEN\_STRIDED\_XFER**.

For transfers of type **LAPI\_GEN\_GENERIC** and **LAPI\_GEN\_IOVECTOR**, the fields are used as follows:

*num\_vecs* indicates the number of data vectors to transfer. Each data vector is defined by a base address and data length.

*info* is the array of addresses.

*len* is the array of data lengths.

For example, consider the following vector description:

```
vec_type = LAPI_GEN_IOVECTOR
num_vecs = 3
info     = {addr_0, addr_1, addr_2}
len      = {len_0, len_1, len_2}
```

On the origin side, this example would tell LAPI to read *len\_0* bytes from *addr\_0*, *len\_1* bytes from *addr\_1*, and *len\_2* bytes from *addr\_2*. As a target vector, this example would tell LAPI to write *len\_0* bytes to *addr\_0*, *len\_1* bytes to *addr\_1*, and *len\_2* bytes to *addr\_2*.

Recall that vector transfers require an origin and target vector. For **LAPI\_Amsendv** calls, the origin vector is passed to the API call on the origin task. The address of the target vector is returned by the header handler.

For transfers of type **LAPI\_GEN\_GENERIC**, the target vector description must also have type **LAPI\_GEN\_GENERIC**. The contents of the *info* and *len* arrays are unrestricted in the generic case; the number of vectors and the length of vectors on the origin and target do not need to match. In this case, LAPI transfers a given number of bytes in noncontiguous buffers specified by the origin vector to a set of noncontiguous buffers specified by the target vector.

If the sum of target vector data lengths (say **TGT\_LEN**) is less than the sum of origin vector data lengths (say **ORG\_LEN**), only the first **TGT\_LEN** bytes from the origin buffers are transferred and the remaining bytes are discarded. If **TGT\_LEN** is greater than **ORG\_LEN**, all **ORG\_LEN** bytes are transferred. Consider the following example:

```
Origin_vector: {
    num_vecs = 3;
    info     = {orgaddr_0, orgaddr_1, orgaddr_2};
    len      = {5, 10, 5}
}

Target_vector: {
    num_vecs = 4;
    info     = {tgtaddr_0, tgtaddr_1, tgtaddr_2, tgtaddr_3};
    len      = {12, 2, 4, 2}
}
```

LAPI copies data as follows:

1. 5 bytes from orgaddr\_0 to tgtaddr\_0 (leaves 7 bytes of space at a 5-byte offset from tgtaddr\_0)
2. 7 bytes from orgaddr\_1 to remaining space in tgtaddr\_0 (leaves 3 bytes of data to transfer from orgaddr\_1)
3. 2 bytes from orgaddr\_1 to tgtaddr\_1 (leaves 1 byte to transfer from orgaddr\_1)
4. 1 byte from orgaddr\_1 followed by 3 bytes from orgaddr\_2 to tgt\_addr\_2 (leaves 3 bytes to transfer from orgaddr\_2)
5. 2 bytes from orgaddr\_2 to tgtaddr\_3

LAPI will copy data from the origin until the space described by the target is filled. For example:

```
Origin_vector: {
    num_vecs = 1;
    info     = {orgaddr_0};
    len      = {20}
}

Target_vector: {
    num_vecs = 2;
    info     = {tgtaddr_0, tgtaddr_1};
    len      = {5, 10}
}
```

LAPI will copy 5 bytes from orgaddr\_0 to tgtaddr\_0 and the next 10 bytes from orgaddr\_0 to tgtaddr\_1. The remaining 5 bytes from orgaddr\_0 will not be copied.

For transfers of type **LAPI\_GEN\_IOVECTOR**, the lengths of the vectors must match and the target vector description must match the origin vector description. More specifically, the target vector description must:

- also have type **LAPI\_GEN\_IOVECTOR**
- have the same *num\_vecs* as the origin vector
- initialize the *info* array with *num\_vecs* addresses in the target address space. For LAPI vectors origin\_vector and target\_vector described similarly to the example above, data is copied as follows:

1. transfer `origin_vector.len[0]` bytes from the address at `origin_vector.info[0]` to the address at `target_vector.info[0]`
2. transfer `origin_vector.len[1]` bytes from the address at `origin_vector.info[1]` to the address at `target_vector.info[1]`
3. transfer `origin_vector.len[n]` bytes from the address at `origin_vector.info[n]` to the address at `target_vector.info[n]`, for  $n = 2$  to  $n = [num\_vecs-3]$
4. transfer `origin_vector.len[num_vecs-2]` bytes from the address at `origin_vector.info[num_vecs-2]` to the address at `target_vector.info[num_vecs-2]`
5. copy `origin_vector.len[num_vecs-1]` bytes from the address at `origin_vector.info[num_vecs-1]` to the address at `target_vector.info[num_vecs-1]`

### Strided vector transfers

For transfers of type **LAPI\_GEN\_STRIDED\_XFER**, the target vector description must match the origin vector description. Rather than specifying the set of address and length pairs, the *info* array of the origin and target vectors is used to specify a data block "template", consisting of a base address, block size and stride. LAPI thus expects the *info* array to contain three integers. The first integer contains the base address, the second integer contains the block size to copy, and the third integer contains the byte stride. In this case, *num\_vecs* indicates the number of blocks of data that LAPI should copy, where the first block begins at the base address. The number of bytes to copy in each block is given by the block size and the starting address for all but the first block is given by previous address + stride. The total amount of data to be copied will be  $num\_vecs * block\_size$ . Consider the following example:

```
Origin_vector {
    num_vecs = 3;
    info     = {orgaddr, 5, 8}
}
```

Based on this description, LAPI will transfer 5 bytes from `orgaddr`, 5 bytes from `orgaddr+8` and 5 bytes from `orgaddr+16`.

### Call details

As mentioned above, counter and handler behavior in **LAPI\_Amsendv** is nearly identical to that of **LAPI\_Amsend**. A short summary of that behavior is provided here. See the **LAPI\_Amsend** description for full details.

This is a non-blocking call. The calling task cannot change the *uhdr* (origin header) and *org\_vec* data until completion at the origin is signaled by the *org\_cntr* being incremented. The calling task cannot assume that the *org\_vec* structure can be changed before the origin counter is incremented. The structure (of type **lapi\_vec\_t**) that is returned by the header handler cannot be modified before the target counter has been incremented. Also, if a completion handler is specified, it may execute asynchronously, and can only be assumed to have completed after the target counter increments (on the target) or the completion counter increments (at the origin).

The length of the user-specified header (*uhdr\_len*) is constrained by the implementation-specified maximum value **MAX\_UHDR\_SZ**. *uhdr\_len* must be a multiple of the processor's doubleword size. To get the best bandwidth, *uhdr\_len* should be as small as possible.

If the following requirement is not met, an error condition occurs:

- If a strided vector is being transferred, the size of each block must not be greater than the stride size in bytes.

LAPI does not check for any overlapping regions among vectors either at the origin or the target. If the overlapping regions exist on the target side, the contents of the target buffer are undefined after the operation.

## Parameters

<i>hdl</i>	Specifies the LAPI handle.
<i>tgt</i>	Specifies the task ID of the target task. The value of this parameter must be in the range $0 \leq \text{tgt} < \text{NUM\_TASKS}$ .
<i>hdr_hdl</i>	Points to the remote header handler function to be invoked at the target. The value of this parameter can take an address handle that had been previously registered using the <b>LAPI_Addr_set/LAPI_Addr_get</b> mechanism. The value of this parameter cannot be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<i>uhdr</i>	Specifies the pointer to the local header (parameter list) that is passed to the handler function. If <i>uhdr_len</i> is 0, The value of this parameter can be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<i>uhdr_len</i>	Specifies the length of the user's header. The value of this parameter must be a multiple of the processor's doubleword size in the range $0 \leq \text{uhdr\_len} \leq \text{MAX\_UHDR\_SZ}$ .
<i>org_vec</i>	Points to the origin vector.

## INPUT/OUTPUT

<i>tgt_cntr</i>	Specifies the target counter address. The target counter is incremented after the completion handler (if specified) completes or after the completion of data transfer. If the value of this parameter is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), the target counter is not updated.
<i>org_cntr</i>	Specifies the origin counter address (in C) or the origin counter (in FORTRAN). The origin counter is incremented after data is copied out of the origin address (in C) or the origin (in FORTRAN). If the value of this parameter is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), the origin counter is not updated.
<i>cmpl_cntr</i>	Specifies the counter at the origin that signifies completion of the completion handler. It is updated once the completion handler completes. If no completion handler is specified, the counter is incremented at the completion of message delivery. If the value of this parameter is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), the completion counter is not updated.

## OUTPUT

<i>ierror</i>	Specifies a FORTRAN return code. This is always the last parameter.
---------------	---

## C Examples

1. To send a **LAPI\_GEN\_IOVECTOR** using active messages:

```
/* header handler routine to execute on target task */
lapi_vec_t *hdr_hdlr(lapi_handle_t *handle, void *uhdr, uint *uhdr_len,
                    ulong *len_vec[ ], compl_hdlr_t **completion_handler,
                    void **user_info)
{
    /* set completion handler pointer and other info */

    /* set up the vector to return to LAPI */
    /* for a LAPI_GEN_IOVECTOR: num_vecs, vec_type, and len must all have */
    /* the same values as the origin vector. The info array should */
    /* contain the buffer addresses for LAPI to write the data */
    vec->num_vecs = NUM_VECS;
    vec->vec_type = LAPI_GEN_IOVECTOR;
    vec->len = (unsigned long *)malloc(NUM_VECS*sizeof(unsigned long));
    vec->info = (void **) malloc(NUM_VECS*sizeof(void *));
    for( i=0; i < NUM_VECS; i++ ) {
        vec->info[i] = (void *) &data_buffer[i];
        vec->len[i] = (unsigned long)(sizeof(int));
    }
}
```

```

    }

    return vec;
}

{
    .
    .
    .

    void      *hdr_hdlr_list[NUM_TASKS]; /* table of remote header handlers */
    lapi_vec_t *vec;                    /* data for data transfer */

    vec->num_vecs = NUM_VECS;
    vec->vec_type = LAPI_GEN_IOVECTOR;
    vec->len      = (unsigned long *) malloc(NUM_VECS*sizeof(unsigned long));
    vec->info     = (void **) malloc(NUM_VECS*sizeof(void *));

    /* each vec->info[i] gets a base address */
    /* each vec->len[i] gets the number of bytes to transfer from vec->info[i] */

    LAPI_Amsendv(hdr_hdl, tgt, (void *) hdr_hdl_list[buddy], NULL, 0, vec,
                 tgt_cntr, org_cntr, compl_cntr);

    /* data will be copied as follows: */
    /* len[0] bytes of data starting from address info[0] */
    /* len[1] bytes of data starting from address info[1] */
    .
    .
    .
    /* len[NUM_VECS-1] bytes of data starting from address info[NUM_VECS-1] */
}

```

The above example could also illustrate the **LAPI\_GEN\_GENERIC** type, with the following modifications:

- Both vectors would need **LAPI\_GEN\_GENERIC** as the `vec_type`.
- There are no restrictions on symmetry of number of vectors and lengths between the origin and target sides.

## 2. To send a **LAPI\_STRIDED\_VECTOR** using active messages:

```

/* header handler routine to execute on target task */
lapi_vec_t *hdr_hdlr(lapi_handle_t *handle, void *uhdr, uint *uhdr_len,
                    ulong *len_vec[ ], compl_hdlr_t **completion_handler,
                    void **user_info)
{
    int block_size;          /* block size */
    int data_size;          /* stride */
    .
    .
    .
    vec->num_vecs = NUM_VECS; /* NUM_VECS = number of vectors to transfer */
                          /* must match that of the origin vector */
    vec->vec_type = LAPI_GEN_STRIDED_XFER; /* same as origin vector */

    /* see comments in origin vector setup for a description of how data */
    /* will be copied based on these settings. */
    vec->info[0] = buffer_address; /* starting address for data copy */
    vec->info[1] = block_size;     /* bytes of data to copy */
    vec->info[2] = stride;         /* distance from copy block to copy block */
    .
}

```

```

    .
    .
    return vec;
}

{
    .
    .
    .
    lapi_vec_t *vec;                               /* data for data transfer */

    vec->num_vecs = NUM_VECS;                       /* NUM_VECS = number of vectors to transfer */
                                                    /* must match that of the target vector */
    vec->vec_type = LAPI_GEN_STRIDED_XFER;          /* same as target vector */

    vec->info[0] = buffer_address;                   /* starting address for data copy */
    vec->info[1] = block_size;                       /* bytes of data to copy */
    vec->info[2] = stride;                           /* distance from copy block to copy block */
    /* data will be copied as follows: */
    /* block_size bytes will be copied from buffer_address */
    /* block_size bytes will be copied from buffer_address+stride */
    /* block_size bytes will be copied from buffer_address+(2*stride) */
    /* block_size bytes will be copied from buffer_address+(3*stride) */
    .
    .
    .
    /* block_size bytes will be copied from buffer_address+((NUM_VECS-1)*stride) */
    .
    .
    .
    /* if uhdr isn't used, uhdr should be NULL and uhdr_len should be 0 */
    /* tgt_cntr, org_cntr and cpl_cntr can all be NULL */
    LAPI_Amsendv(hndl, tgt, (void *) hdr_hdl_list[buddy], uhdr, uhdr_len,
                 vec, tgt_cntr, org_cntr, cpl_cntr);
    .
    .
}

```

For complete examples, see the sample programs shipped with LAPI.

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HDR_HNDLR_NULL</b>	Indicates that the <i>hdr_hdl</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_ORG_EXTENT</b>	Indicates that the <i>org_vec</i> 's extent ( $\text{stride} * \text{num\_vecs}$ ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_ORG_STRIDE</b>	Indicates that the <i>org_vec</i> stride is less than block.
<b>LAPI_ERR_ORG_VEC_ADDR</b>	Indicates that the <i>org_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but its length ( <i>org_vec-&gt;len[i]</i> ) is not 0.
<b>LAPI_ERR_ORG_VEC_LEN</b>	Indicates that the sum of <i>org_vec-&gt;len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .



<b>LAPI_ERR_ORG_VEC_NULL</b>	Indicates that <i>org_vec</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_ORG_VEC_TYPE</b>	Indicates that the <i>org_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_STRIDE_ORG_VEC_ADDR_NULL</b>	Indicates that the strided vector address <i>org_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.
<b>LAPI_ERR_UHDR_LEN</b>	Indicates that the <i>uhdr_len</i> value passed in is greater than <b>MAX_UHDR_SZ</b> or is not a multiple of the processor's doubleword size.
<b>LAPI_ERR_UHDR_NULL</b>	Indicates that the <i>uhdr</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but <i>uhdr_len</i> is not 0.

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address\_init**, **LAPI\_Addr\_get**, **LAPI\_Addr\_set**, **LAPI\_Amsend**, **LAPI\_Getcntr**, **LAPI\_Getv**, **LAPI\_Putv**, **LAPI\_Qenv**, **LAPI\_Waitcntr**, **LAPI\_Xfer**

---

## LAPI\_Fence Subroutine

### Purpose

Enforces order on LAPI calls.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Fence(hdl)
lapi_handle_t hdl;
```

### FORTRAN Syntax

```
include 'lapif.h'

LAPI_FENCE(hdl, ierror)
INTEGER hdl
INTEGER ierror
```

### Description

**Type of call:** Local data synchronization (blocking) (may require progress on the remote task)

Use this subroutine to enforce order on LAPI calls. If a task calls **LAPI\_Fence**, all the LAPI operations that were initiated by that task, before the fence using the LAPI context *hdl*, are guaranteed to complete at the target tasks. This occurs before any of its communication operations using *hdl*, initiated after the

**LAPI\_Fence**, start transmission of data. This is a data fence which means that the data movement is complete. This is not an operation fence which would need to include active message completion handlers completing on the target.

**LAPI\_Fence** may require internal protocol processing on the remote side to complete the fence request.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

## C Examples

To establish a data barrier in a single task:

```
lapi_handle_t hndl; /* the LAPI handle */
:
/* API communication call 1 */
/* API communication call 2 */
:
/* API communication call n */
LAPI_Fence(hndl);
/* all data movement from above communication calls has completed by this point */
/* any completion handlers from active message calls could still be running. */
```

## Location

**/usr/lib/liblapi\_r.a**

## Related Information

Subroutines: **LAPI\_Amsend**, **LAPI\_Gfence**

---

## LAPI\_Get Subroutine

### Purpose

Copies data from a remote task to a local task.

### Library

Availability Library (**liblapi\_r.a**)

## C Syntax

```
#include <lapi.h>
```

```
int LAPI_Get(hndl, tgt, len, tgt_addr, org_addr, tgt_cntr, org_cntr)
lapi_handle_t hndl;
uint         tgt;
ulong        len;
void         *tgt_addr;
void         *org_addr;
lapi_cntr_t  *tgt_cntr;
lapi_cntr_t  *org_cntr;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_GET(hndl, tgt, len, tgt_addr, org_addr, tgt_cntr, org_cntr, ierror)
INTEGER hndl
INTEGER tgt
INTEGER (KIND=LAPI_LONG_TYPE) :: len
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_addr
INTEGER (KIND=LAPI_ADDR_TYPE) :: org_addr
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_cntr
TYPE (LAPI_CNTR_T) :: org_cntr
INTEGER ierror
```

## Description

**Type of call:** point-to-point communication (non-blocking)

Use this subroutine to transfer data from a remote (target) task to a local (origin) task. Note that in this case the origin task is actually the *receiver* of the data. This difference in transfer type makes the counter behavior slightly different than in the normal case of origin sending to target.

The origin buffer will still increment on the origin task upon availability of the origin buffer. But in this case, the origin buffer becomes available once the transfer of data is complete. Similarly, the target counter will increment once the target buffer is available. Target buffer availability in this case refers to LAPI no longer needing to access the data in the buffer.

This is a non-blocking call. The caller *cannot* assume that data transfer has completed upon the return of the function. Instead, counters should be used to ensure correct buffer addresses as defined above.

Note that a zero-byte message does not transfer data, but it does have the same semantic with respect to counters as that of any other message.

## Parameters

### INPUT

<i>hndl</i>	Specifies the LAPI handle.
<i>tgt</i>	Specifies the task ID of the target task that is the source of the data. The value of this parameter must be in the range $0 \leq tgt < \text{NUM\_TASKS}$ .
<i>len</i>	Specifies the number of bytes of data to be copied. This parameter must be in the range $0 \leq len \leq$ the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<i>tgt_addr</i>	Specifies the target buffer address of the data source. If <i>len</i> is <b>0</b> , The value of this parameter can be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).

### INPUT/OUTPUT

<i>tgt_cntr</i>	Specifies the target counter address. The target counter is incremented once the data
-----------------	---

buffer on the target can be modified. If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the target counter is not updated.

*org\_cntr* Specifies the origin counter address (in C) or the origin counter (in FORTRAN). The origin counter is incremented after data arrives at the origin. If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the origin counter is not updated.

## OUTPUT

*org\_addr* Specifies the local buffer address into which the received data is copied. If *len* is 0, The value of this parameter can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

```
{
  /* initialize the table buffer for the data addresses */

  /* get remote data buffer addresses */
  LAPI_Address_init(hndl, (void *)data_buffer, data_buffer_list);
  .
  .
  LAPI_Get(hndl, tgt, (ulong) data_len, (void *) (data_buffer_list[tgt]),
           (void *) data_buffer, tgt_cntr, org_cntr);

  /* retrieve data_len bytes from address data_buffer_list[tgt] on task tgt. */
  /* write the data starting at address data_buffer. tgt_cntr and org_cntr */
  /* can be NULL. */
}
```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_DATA_LEN</b>	Indicates that the value of <i>udata_len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_ORG_ADDR_NULL</b>	Indicates that the <i>org_addr</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but <i>len</i> is greater than 0.
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_ADDR_NULL</b>	Indicates that the <i>tgt_addr</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but <i>len</i> is greater than 0.
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address\_init**, **LAPI\_Getcntr**, **LAPI\_Put**, **LAPI\_Qenv**, **LAPI\_Waitcntr**, **LAPI\_Xfer**

---

## LAPI\_Getcntr Subroutine

### Purpose

Gets the integer value of a specified LAPI counter.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Getcntr(hdl, cntr, val)
lapi_handle_t hdl;
lapi_cntr_t *cntr;
int *val;
```

### FORTRAN Syntax

```
include 'lapif.h'

LAPI_GETCNTR(hdl, cntr, val, ierror)
INTEGER hdl
TYPE (LAPI_CNTR_T) :: cntr
INTEGER val
INTEGER ierror
```

### Description

**Type of call:** Local counter manipulation

This subroutine gets the integer value of *cntr*. It is used to check progress on *hdl*.

### Parameters

#### INPUT

*hdl* Specifies the LAPI handle.

*cntr* Specifies the address of the counter. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

#### OUTPUT

*val* Returns the integer value of the counter *cntr*. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

### C Examples

```
{
    lapi_cntr_t cntr;
    int val;

    /* cntr is initialized */

    /* processing/communication takes place */
```

```

LAPI_Getcctr(hndl, &cntr, &val)

/* val now contains the current value of cntr */
}

```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_CNTR\_NULL** Indicates that the *cntr* pointer is NULL (in C) or that the value of *cntr* is **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_RET\_PTR\_NULL** Indicates that the value of the *val* pointer is NULL (in C) or that the value of *val* is **LAPI\_ADDR\_NULL** (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Amsend**, **LAPI\_Amsendv**, **LAPI\_Get**, **LAPI\_Getv**, **LAPI\_Put**, **LAPI\_Putv**, **LAPI\_Rmw**, **LAPI\_Setcctr**, **LAPI\_Waitcctr**, **LAPI\_Xfer**

---

## LAPI\_Getv Subroutine

### Purpose

Copies vectors of data from a remote task to a local task.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```

#include <lapi.h>

int LAPI_Getv(hndl, tgt, tgt_vec, org_vec, tgt_cntr, org_cntr)
lapi_handle_t hndl;
uint          tgt;
lapi_vec_t   *tgt_vec;
lapi_vec_t   *org_vec;
lapi_cntr_t  *tgt_cntr;
lapi_cntr_t  *org_cntr;

typedef struct {
    lapi_vectype_t  vec_type; /* operation code */
    uint           num_vecs; /* number of vectors */
    void           **info;   /* vector of information */
    ulong          *len;     /* vector of lengths */
} lapi_vec_t;

```

### FORTRAN Syntax

```

include 'lapif.h'

LAPI_GETV(hndl, tgt, tgt_vec, org_vec, tgt_cntr, org_cntr, ierror)
INTEGER hndl
INTEGER tgt

```

```

INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_vec
TYPE (LAPI_VEC_T) :: org_vec
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_cntr
TYPE (LAPI_CNTR_T) :: org_cntr
INTEGER ierror

```

The 32-bit version of the **LAPI\_VEC\_T** type is defined as:

```

TYPE LAPI_VEC_T
  SEQUENCE
    INTEGER(KIND = 4) :: vec_type
    INTEGER(KIND = 4) :: num_vecs
    INTEGER(KIND = 4) :: info
    INTEGER(KIND = 4) :: len
END TYPE LAPI_VEC_T

```

The 64-bit version of the **LAPI\_VEC\_T** type is defined as:

```

TYPE LAPI_VEC_T
  SEQUENCE
    INTEGER(KIND = 4) :: vec_type
    INTEGER(KIND = 4) :: num_vecs
    INTEGER(KIND = 8) :: info
    INTEGER(KIND = 8) :: len
END TYPE LAPI_VEC_T

```

## Description

**Type of call:** point-to-point communication (non-blocking)

This subroutine is the vector version of the **LAPI\_Get** call. Use **LAPI\_Getv** to transfer vectors of data from the target task to the origin task. Both the origin and target vector descriptions are located in the address space of the origin task. But, the values specified in the *info* array of the target vector must be addresses in the address space of the target task.

The calling program *cannot* assume that the origin buffer can be changed or that the contents of the origin buffers on the origin task are ready for use upon function return. After the origin counter (*org\_cntr*) is incremented, the origin buffers can be modified by the origin task. After the target counter (*tgt\_cntr*) is incremented, the target buffers can be modified by the target task. If you provide a completion counter (*cmpl\_cntr*), it is incremented at the origin after the target counter (*tgt\_cntr*) has been incremented at the target. If the values of any of the counters or counter addresses are NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the data transfer occurs, but the corresponding counter increments do not occur.

If any of the following requirements are not met, an error condition occurs:

- The vector types *org\_vec*→*vec\_type* and *tgt\_vec*→*vec\_type* must be the same.
- If a strided vector is being transferred, the size of each block must not be greater than the stride size in bytes.
- The length of any vector that is pointed to by *tgt\_vec* must be equal to the length of the corresponding vector that is pointed to by *org\_vec*.

LAPI does not check for any overlapping regions among vectors either at the origin or the target. If the overlapping regions exist on the origin side, the contents of the origin buffer are undefined after the operation.

See **LAPI\_Amsendv** for details about communication using different LAPI vector types. (**LAPI\_Getv** does not support the **LAPI\_GEN\_GENERIC** type.)

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*tgt* Specifies the task ID of the target task. The value of this parameter must be in the range 0 <= *tgt* < **NUM\_TASKS**.

*tgt\_vec* Points to the target vector description.

*org\_vec* Points to the origin vector description.

#### INPUT/OUTPUT

*tgt\_cntr* Specifies the target counter address. The target counter is incremented once the data buffer on the target can be modified. If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the target counter is not updated.

*org\_cntr* Specifies the origin counter address (in C) or the origin counter (in FORTRAN). The origin counter is incremented after data arrives at the origin. If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the origin counter is not updated.

#### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To get a LAPI\_GEN\_IOVECTOR:

```
{
    /* retrieve a remote data buffer address for data to transfer, */
    /* such as through LAPI_Address_init */

    /* task that calls LAPI_Getv sets up both org_vec and tgt_vec */
    org_vec->num_vecs = NUM_VECS;
    org_vec->vec_type = LAPI_GEN_IOVECTOR;
    org_vec->len = (unsigned long *)
    malloc(NUM_VECS*sizeof(unsigned long));
    org_vec->info = (void **) malloc(NUM_VECS*sizeof(void *));

    /* each org_vec->info[i] gets a base address on the origin task */
    /* each org_vec->len[i] gets the number of bytes to write to */
    /* org_vec->info[i] */

    tgt_vec->num_vecs = NUM_VECS;
    tgt_vec->vec_type = LAPI_GEN_IOVECTOR;
    tgt_vec->len = (unsigned long *)
    malloc(NUM_VECS*sizeof(unsigned long));
    tgt_vec->info = (void **) malloc(NUM_VECS*sizeof(void *));

    /* each tgt_vec->info[i] gets a base address on the target task */
    /* each tgt_vec->len[i] gets the number of bytes to transfer */
    /* from vec->info[i] */
    /* For LAPI_GEN_IOVECTOR, num_vecs, vec_type, and len must be */
    /* the same */

    LAPI_Getv(hndl, tgt, tgt_vec, org_vec, tgt_cntr, org_cntr);
    /* tgt_cntr and org_cntr can both be NULL */

    /* data will be retrieved as follows: */
    /* org_vec->len[0] bytes will be retrieved from */
    /* tgt_vec->info[0] and written to org_vec->info[0] */
    /* org_vec->len[1] bytes will be retrieved from */
    /* tgt_vec->info[1] and written to org_vec->info[1] */
    .
    .
    .
    /* org_vec->len[NUM_VECS-1] bytes will be retrieved */
}
```



```

    /* from tgt_vec->info[NUM_VECS-1] and written to */
    /* org_vec->info[NUM_VECS-1] */
}

```

For examples of other vector types, see **LAPI\_Amsendv**.

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_ORG_EXTENT</b>	Indicates that the <i>org_vec</i> 's extent (stride * <i>num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_ORG_STRIDE</b>	Indicates that the <i>org_vec</i> stride is less than block size.
<b>LAPI_ERR_ORG_VEC_ADDR</b>	Indicates that some <i>org_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but the corresponding length ( <i>org_vec-&gt;len[i]</i> ) is not 0.
<b>LAPI_ERR_ORG_VEC_LEN</b>	Indicates that the total sum of all <i>org_vec-&gt;len[i]</i> (where [i] is in the range 0 <= i <= <i>org_vec-&gt;num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_ORG_VEC_NULL</b>	Indicates that the <i>org_vec</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_ORG_VEC_TYPE</b>	Indicates that the <i>org_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_STRIDE_ORG_VEC_ADDR_NULL</b>	Indicates that the strided vector base address <i>org_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_STRIDE_TGT_VEC_ADDR_NULL</b>	Indicates that the strided vector address <i>tgt_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_EXTENT</b>	Indicates that <i>tgt_vec</i> 's extent (stride * <i>num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.
<b>LAPI_ERR_TGT_STRIDE</b>	Indicates that the <i>tgt_vec</i> 's stride is less than its block size.
<b>LAPI_ERR_TGT_VEC_ADDR</b>	Indicates that the <i>tgt_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but its length ( <i>tgt_vec-&gt;len[i]</i> ) is not 0.
<b>LAPI_ERR_TGT_VEC_LEN</b>	Indicates that the sum of <i>tgt_vec-&gt;len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_TGT_VEC_NULL</b>	Indicates that <i>tgt_vec</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT_VEC_TYPE</b>	Indicates that the <i>tgt_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_VEC_LEN_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different lengths ( <i>len[]</i> ).
<b>LAPI_ERR_VEC_NUM_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different <i>num_vecs</i> .

**LAPI\_ERR\_VEC\_TYPE\_DIFF**

Indicates that *org\_vec* and *tgt\_vec* have different vector types (*vec\_type*).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Amsendv**, **LAPI\_Getcntr**, **LAPI\_Putv**, **LAPI\_Qenv**, **LAPI\_Waitcntr**

---

## LAPI\_Gfence Subroutine

### Purpose

Enforces order on LAPI calls across all tasks and provides barrier synchronization among them.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Gfence(hdl)
lapi_handle_t hdl;
```

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_GFENCE(hdl, ierror)
INTEGER hdl
INTEGER ierror
```

### Description

**Type of call:** collective data synchronization (blocking)

Use this subroutine to enforce global order on LAPI calls. This is a *collective call*. Collective calls must be made in the same order at all participating tasks.

On completion of this call, it is assumed that all LAPI communication associated with *hdl* from all tasks has quiesced. Although *hdl* is local, it represents a set of tasks that were associated with it at **LAPI\_Init**, all of which must participate in this operation for it to complete. This is a data fence, which means that the data movement is complete. This is not an operation fence, which would need to include active message completion handlers completing on the target.

### Parameters

#### INPUT

*hdl* Specifies the LAPI handle.

#### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Fence**

---

## LAPI\_Init Subroutine

### Purpose

Initializes a LAPI context.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Init(hndl,lapi_info)
lapi_handle_t *hndl;
lapi_info_t *lapi_info;
```

### FORTRAN Syntax

```
include 'lapif.h'

LAPI_INIT(hndl,lapi_info,ierror)
INTEGER hndl
TYPE (LAPI_INFO_T) :: lapi_info
INTEGER ierror
```

### Description

**Type of call:** Local initialization

Use this subroutine to instantiate and initialize a new LAPI context. A handle to the newly-created LAPI context is returned in *hndl*. All subsequent LAPI calls can use *hndl* to specify the context of the LAPI operation. Except for **LAPI\_Address()** and **LAPI\_Msg\_string()**, the user cannot make any LAPI calls before calling **LAPI\_Init()**.

The *lapi\_info* structure (**lapi\_info\_t**) must be "zeroed out" before any fields are filled in. To do this in C, use this statement: **bzero (lapi\_info, size of (lapi\_info\_t))**. In FORTRAN, you need to "zero out" each field manually in the **LAPI\_INFO\_T** type. Fields with a description of Future support should not be used because the names of those fields might change.

The **lapi\_info\_t** structure is defined as follows:

```
typedef struct {
    lapi_dev_t    protocol;        /* Protocol device returned          */
    lapi_lib_t    lib_vers;       /* LAPI library version -- user-supplied */
    uint         epoch_num;      /* No longer used                    */
}
```

```

int          num_compl_hdlr_thr; /* Number of completion handler threads */
uint        instance_no; /* Instance of LAPI to initialize [1-16] */
int         info6; /* Future support */
LAPI_err_hdlr *err_hdlr; /* User-registered error handler */
com_thread_info_t *lapi_thread_attr; /* Support thread attr and init function */
void        *adapter_name; /* What adapter to initialize, i.e. css0, m10 */
lapi_extnd_t *add_info; /* Additional structure extension */
} lapi_info_t;

```

The fields are used as follows:

*protocol* LAPI sets this field to the protocol that has been initialized.

*lib\_vers* Is used to indicate a library version to LAPI for compatibility purposes. Valid values for this field are:

- L1\_LIB** Provides basic functionality (this is the default).
- L2\_LIB** Provides the ability to use counters as structures.
- LAST\_LIB** Provides the most current level of functionality. For new users of LAPI, *lib\_vers* should be set to **LAST\_LIB**.

This field must be set to **L2\_LIB** or **LAST\_LIB** to use **LAPI\_Nopoll\_wait** and **LAPI\_Setcntr\_wstatus**.

*epoch\_num* This field is no longer used.

*num\_compl\_hdlr\_thr* Indicates to LAPI the number of completion handler threads to initialize.

*instance\_no* Specifies the instance of LAPI to initialize (**1** to **16**).

*info6* This field is for future use.

*err\_hdlr* Use this field to optionally pass a callback pointer to an error-handler routine.

*lapi\_thread\_attr* Supports thread attributes and initialization function.

*adapter\_name* Is used in persistent subsystem (PSS) mode to pass an adapter name.

*add\_info* Is used for additional information in standalone UDP mode.

## Parameters

### INPUT/OUTPUT

*lapi\_info* Specifies a structure that provides the parallel job information with which this LAPI context is associated. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

### OUTPUT

*hdl* Specifies a pointer to the LAPI handle to initialize.

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_ALL\_HNDL\_IN\_USE** All available LAPI instances are in use.

**LAPI\_ERR\_BOTH\_NETSTR\_SET**

Both the **MP\_LAPI\_NETWORK** and **MP\_LAPI\_INET** statements are set (only one should be set).

**LAPI\_ERR\_CSS\_LOAD\_FAILED**

LAPI is unable to load the communication utility library.

**LAPI\_ERR\_HNDL\_INVALID**

The **lapi\_handle\_t** \* passed to LAPI for initialization is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_INFO\_NONZERO\_INFO**

The future support fields in the **lapi\_info\_t** structure that was passed to LAPI are not set to zero (and should be).

**LAPI\_ERR\_INFO\_NULL**

The **lapi\_info\_t** pointer passed to LAPI is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_MEMORY\_EXHAUSTED**

LAPI is unable to obtain memory from the system.

**LAPI\_ERR\_MSG\_API**

Indicates that the **MP\_MSG\_API** environment variable is not set correctly.

**LAPI\_ERR\_NO\_NETSTR\_SET**

No network statement is set. Note that if running with POE, this will be returned if **MP\_MSG\_API** is not set correctly.

**LAPI\_ERR\_NO\_UDP\_HNDLR**

You passed a value of NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN) for both the UDP handler and the UDP list. One of these (the UDP handler or the UDP list) must be initialized for standalone UDP initialization. This error is returned in standalone UDP mode only.

**LAPI\_ERR\_PSS\_NON\_ROOT**

You tried to initialize the persistent subsystem (PSS) protocol as a non-**root** user.

**LAPI\_ERR\_SHM\_KE\_NOT\_LOADED**

LAPI's shared memory kernel extension is not loaded.

**LAPI\_ERR\_SHM\_SETUP**

LAPI is unable to set up shared memory. This error will be returned if **LAPI\_USE\_SHM=only** and tasks are assigned to more than one node.

**LAPI\_ERR\_UDP\_PKT\_SZ**

The UDP packet size you indicated is not valid.

**LAPI\_ERR\_UNKNOWN**

An internal error has occurred.

**LAPI\_ERR\_USER\_UDP\_HNDLR\_FAIL**

The UDP handler you passed has returned a non-zero error code. This error is returned in standalone UDP mode only.

## C Examples

The following environment variable must be set before LAPI is initialized:

```
MP_MSG_API=[ lapi | [ lapi,mpi | mpi,lapi ] | mpi_lapi ]
```

The following environment variables are also commonly used:

```
MP_EUILIB=[ ip | us ] (ip is the default)
```

```
MP_PROCS=number_of_tasks_in_job
```

```
LAPI_USE_SHM=[ yes | no | only ] (no is the default)
```

To initialize LAPI, follow these steps:

1. Set environment variables (as described in *RSCT for AIX 5L: LAPI Programming Guide*) before the user application is invoked. The remaining steps are done in the user application.
2. Clear `lapi_info_t`, then set any fields.
3. Call `LAPI_Init`.

### For systems running PE

Both US and UDP/IP are supported for shared handles as long as they are the same for both handles. Mixed transport protocols such as LAPI IP and shared user space (US) are not supported.

To initialize a LAPI handle:

```
{
    lapi_handle_t hndl;
    lapi_info_t info;

    bzero(&info, sizeof(lapi_info_t)); /* clear lapi_info */

    LAPI_Init(&hndl, &info);
}
```

To initialize a LAPI handle and register an error handler:

```
void my_err_hdlr(lapi_handle_t *hndl, int *error_code, lapi_err_t *err_type,
                int *task_id, int *src )
{
    /* examine passed parameters and delete desired information */

    if ( user wants to terminate ) {
        LAPI_Term(*hndl);          /* will terminate LAPI */
        exit(some_return_code);
    }

    /* any additional processing */

    return; /* signals to LAPI that error is non-fatal; execution should continue */
}

{
    lapi_handle_t hndl;
    lapi_info_t info;

    bzero(&info, sizeof(lapi_info_t)); /* clear lapi_info */

    /* set error handler pointer */
    info.err_hdlr = (LAPI_err_hdlr) my_err_hdlr;

    LAPI_Init(&hndl, &info);
}
```

### For standalone systems (not running PE)

To initialize a LAPI handle for UDP/IP communication using a user handler:

```
int my_udp_hdlr(lapi_handle_t *hndl, lapi_udp_t *local_addr, lapi_udp_t *addr_list,
               lapi_udpinfo_t *info)
{
    /* LAPI will allocate and free addr_list pointer when using */
    /* a user handler */
}
```

```

/* use the AIX inet_addr call to convert an IP address      */
/* from a dotted quad to a long                            */
task_0_ip_as_long = inet_addr(task_0_ip_as_string);
addr_list[0].ip_addr = task_0_ip_as_long;
addr_list[0].port_no = task_0_port_as_unsigned;

task_1_ip_as_long = inet_addr(task_1_ip_as_string);
addr_list[1].ip_addr = task_1_ip_as_long;
addr_list[1].port_no = task_1_port_as_unsigned;
.
.
.
task_num_tasks-1_ip_as_long = inet_addr(task_num_tasks-1_ip_as_string);
addr_list[num_tasks-1].ip_addr = task_num_tasks-1_ip_as_long;
addr_list[num_tasks-1].port_no = task_num_tasks-1_port_as_unsigned;
}

{
lapi_handle_t hndl;
lapi_info_t info;
lapi_extend_t extend_info;

bzero(&info, sizeof(lapi_info_t));          /* clear lapi_info      */
bzero(&extend_info, sizeof(lapi_extend_t)); /* clear lapi_extend_info */

extend_info.udp_hdlr = (udp_init_hdlr *) my_udp_hdlr;
info.add_info = &extend_info;

LAPI_Init(&hndl, &info);
}

```

To initialize a LAPI handle for UDP/IP communication using a user list:

```

{
lapi_handle_t hndl;
lapi_info_t info;
lapi_extend_t extend_info;
lapi_udp_t *addr_list;

bzero(&info, sizeof(lapi_info_t));          /* clear lapi_info      */
bzero(&extend_info, sizeof(lapi_extend_t)); /* clear lapi_extend_info */

/* when using a user list, the user is responsible for allocating */
/* and freeing the list pointer                                     */
addr_list = malloc(num_tasks);

/* Note, since we need to know the number of tasks before LAPI is */
/* initialized, we can't use LAPI_Qenv. getenv("MP_PROCS") will   */
/* do the trick.                                                  */

/* populate addr_list                                           */
/* use the AIX inet_addr call to convert an IP address          */
/* from a dotted quad to a long                                  */
task_0_ip_as_long = inet_addr(task_0_ip_as_string);
addr_list[0].ip_addr = task_0_ip_as_long;
addr_list[0].port_no = task_0_port_as_unsigned;

task_1_ip_as_long = inet_addr(task_1_ip_as_string);
addr_list[1].ip_addr = task_1_ip_as_long;
addr_list[1].port_no = task_1_port_as_unsigned;
.
.
.
}

```

```

    .
    task_num_tasks-1_ip_as_long = inet_addr(task_num_tasks-1_ip_as_string);
    addr_list[num_tasks-1].ip_addr = task_num_tasks-1_ip_as_long;
    addr_list[num_tasks-1].port_no = task_num_tasks-1_port_as_unsigned;

    /* then assign to extend pointer */
    extend_info.add_udp_addrs = addr_list;

    info.add_info = &extend_info;

    LAPI_Init(&hdl, &info);
    .
    .
    .

    /* user's responsibility only in the case of user list */
    free(addr_list);
}

```

See the LAPI sample programs for complete examples of initialization in standalone mode.

To initialize a LAPI handle for user space (US) communication in standalone mode:

```

export MP_MSG_API=lapi
export MP_EUILIB=us
export MP_PROCS=          /* number of tasks in job */
export MP_PARTITION=     /* unique job key */
export MP_CHILD=         /* unique task ID */
export MP_LAPI_NETWORK=@1:164,sn0 /* LAPI network information */

run LAPI jobs as normal

```

See the **README.LAPI.STANDALONE.US** file in the **standalone/us** directory of the LAPI sample files for complete details.

## Location

**/usr/lib/liblapi\_r.a**

## Related Information

Books: *RSCT for AIX 5L: LAPI Programming Guide* for information about

- Initializing LAPI on systems running PE
- Initializing LAPI on standalone systems
- Bulk message transfer

Subroutines: **LAPI\_Msg\_string**, **LAPI\_Term**

---

## LAPI\_Msg\_string Subroutine

### Purpose

Retrieves the message that is associated with a subroutine return code.

### Library

Availability Library (**liblapi\_r.a**)



## C Syntax

```
#include <lapi.h>
```

```
LAPI_Msg_string(error_code, buf)
int error_code;
void *buf;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_MSG_STRING(error_code, buf, ierror)
INTEGER error_code
CHARACTER buf(LAPI_MAX_ERR_STRING)
INTEGER ierror
```

## Description

**Type of call:** local queries

Use this subroutine to retrieve the message string that is associated with a LAPI return code. LAPI tries to find the messages of any return codes that come from the AIX operating system or its communication subsystem.

## Parameters

### INPUT

*error\_code* Specifies the return value of a previous LAPI call.

### OUTPUT

*buf* Specifies the buffer to store the message string.

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To get the message string associated with a LAPI return code:

```
{
    char msg_buf[LAPI_MAX_ERR_STRING]; /* constant defined in lapi.h */
    int rc, errc;

    rc = some_LAPI_call();

    errc = LAPI_Msg_string(rc, msg_buf);

    /* msg_buf now contains the message string for the return code */
}
```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_CATALOG\_FAIL** Indicates that the message catalog cannot be opened. An English-only string is copied into the user's message buffer (*buf*).

**LAPI\_ERR\_CODE\_UNKNOWN** Indicates that *error\_code* is outside of the range known to LAPI.

**LAPI\_ERR\_RET\_PTR\_NULL** Indicates that the value of the *buf* pointer is NULL (in C) or that the value of *buf* is **LAPI\_ADDR\_NULL** (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

*R SCT for AIX 5L: LAPI Programming Guide* contains information about

- Initializing LAPI
- Bulk message transfer

Subroutines: **LAPI\_Msg\_string**, **LAPI\_Term**.

---

## LAPI\_Msgpoll Subroutine

### Purpose

Allows the calling thread to check communication progress.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Msgpoll(hndl, cnt, info)  
lapi_handle_t hndl;  
uint cnt;  
lapi_msg_info_t *info;
```

```
typedef struct {  
    lapi_msg_state_t status; /* Message status returned from LAPI_Msgpoll */  
    ulong reserve[10]; /* Reserved */  
} lapi_msg_info_t;
```

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_MSGPOLL(hndl, cnt, info, ierror)  
INTEGER hndl  
INTEGER cnt  
TYPE (LAPI_MSG_STATE_T) :: info  
INTEGER ierror
```

### Description

**Type of call:** local progress monitor (blocking)

The **LAPI\_Msgpoll** subroutine allows the calling thread to check communication progress. With this subroutine, LAPI provides a means of running the dispatcher several times until either progress is made or a specified maximum number of dispatcher loops have executed. Here, *progress* is defined as the completion of either a message send operation or a message receive operation.

**LAPI\_Msgpoll** is intended to be used when interrupts are turned off. If the user has not explicitly turned interrupts off, LAPI temporarily disables interrupt mode while in this subroutine because the dispatcher is called, which will process any pending receive operations. If the LAPI dispatcher loops for the specified

maximum number of times, the call returns. If progress is made before the maximum count, the call will return immediately. In either case, LAPI will report status through a data structure that is passed by reference.

The `lapi_msg_info_t` structure contains a flags field (`status`), which is of type `lapi_msg_state_t`. Flags in the `status` field are set as follows:

<b>LAPI_DISP_CNTR</b>	If the dispatcher has looped <i>cnt</i> times without making progress
<b>LAPI_SEND_COMPLETE</b>	If a message send operation has completed
<b>LAPI_RECV_COMPLETE</b>	If a message receive operation has completed
<b>LAPI_BOTH_COMPLETE</b>	If both a message send operation and a message receive operation have completed
<b>LAPI_POLLING_NET</b>	If another thread is already polling the network or shared memory completion

## Parameters

### INPUT

<i>hdl</i>	Specifies the LAPI handle.
<i>cnt</i>	Specifies the maximum number of times the dispatcher should loop with no progress before returning.
<i>info</i>	Specifies a status structure that contains the result of the <b>LAPI_Msgpoll()</b> call.

### OUTPUT

<i>ierror</i>	Specifies a FORTRAN return code. This is always the last parameter.
---------------	---

## C Examples

To loop through the dispatcher no more than 1000 times, then check what progress has been made:

```
{  
  
    lapi_msg_info_t msg_info;  
    int cnt = 1000;  
    .  
    .  
    .  
    LAPI_Msgpoll(hndl, cnt, &msg_info);  
  
    if ( msg_info.status & LAPI_BOTH_COMPLETE ) {  
        /* both a message receive and a message send have been completed */  
    } else if ( msg_info.status & LAPI_RECV_COMPLETE ) {  
        /* just a message receive has been completed */  
    } else if ( msg_info.status & LAPI_SEND_COMPLETE ) {  
        /* just a message send has been completed */  
    } else {  
        /* cnt loops and no progress */  
    }  
  
}
```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hdl</i> passed in is not valid (not initialized or in terminated state).

## LAPI\_ERR\_MSG\_INFO\_NULL

Indicates that the *info* pointer is NULL (in C) or that the value of *info* is LAPI\_ADDR\_NULL (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: LAPI\_Getcntr, LAPI\_Probe, LAPI\_Setcntr, LAPI\_Waitcntr

---

## LAPI\_Nopoll\_wait Subroutine

### Purpose

Waits for a counter update without polling.

### Library

Availability Library (liblapi\_r.a)

### C Syntax

```
#include <lapi.h>

void LAPI_Nopoll_wait(hndl, cntr_ptr, val, cur_cntr_val)
lapi_handle_t hndl;
lapi_cntr_t *cntr_ptr;
int val;
int *cur_cntr_val;
```

### FORTRAN Syntax

```
include 'lapif.h'

int LAPI_NOPOLL_WAIT(hndl, cntr, val, cur_cntr_val, ierror)
INTEGER hndl
TYPE (LAPI_CNTR_T) :: cntr
INTEGER val
INTEGER cur_cntr_val
INTEGER ierror
```

### Description

**Type of call:** recovery (blocking)

This subroutine waits for a counter update without polling (that is, without explicitly invoking LAPI's internal communication dispatcher). This call may or may not check for message arrivals over the LAPI context *hndl*. The *cur\_cntr\_val* variable is set to the current counter value. Although it has higher latency than LAPI\_Waitcntr, LAPI\_Nopoll\_wait frees up the processor for other uses.

**Note:** To use this subroutine, the *lib\_vers* field in the **lapi\_info\_t** structure must be set to L2\_LIB or LAST\_LIB.

### Parameters

#### INPUT

*hndl* Specifies the LAPI handle.

*val* Specifies the relative counter value (starting from 1) that the counter needs to reach before returning.

*cur\_cntr\_val* Specifies the integer value of the current counter. The value of The value of this parameter can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

## INPUT/OUTPUT

*cntr\_ptr* Points to the **lapi\_cntr\_t** structure in C.

*cntr* Is the **lapi\_cntr\_t** structure in FORTRAN.

## OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_CNTR\_NULL** Indicates that the *cntr\_ptr* pointer is NULL (in C) or that the value of *cntr* is **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_CNTR\_VAL** Indicates that the *val* passed in is less than or equal to 0.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hdl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_MULTIPLE\_WAITERS** Indicates that more than one thread is waiting for the counter.

**LAPI\_ERR\_TGT\_PURGED** Indicates that the subroutine returned early because **LAPI\_Purge\_totask()** was called.

## Restrictions

Use of this subroutine is *not* recommended on a system that is running Parallel Environment (PE).

## Location

`/usr/lib/liblapi_r.a`

## Related Information

Subroutines: **LAPI\_Init**, **LAPI\_Purge\_totask**, **LAPI\_Resume\_totask**, **LAPI\_Setcntr\_wstatus**

---

## LAPI\_Probe Subroutine

### Purpose

Transfers control to the communication subsystem to check for arriving messages and to make progress in polling mode.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Probe(hdl)
lapi_handle_t hdl;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
int LAPI_PROBE(hdl, ierror)
```

```
INTEGER hdl
```

```
INTEGER ierror
```

## Description

**Type of call:** local progress monitor (non-blocking)

This subroutine transfers control to the communication subsystem in order to make progress on messages associated with the context *hdl*. A **LAPI\_Probe** operation lasts for one round of the communication dispatcher.

**Note:** There is no guarantee about receipt of messages on the return from this function.

## Parameters

### INPUT

*hdl* Specifies the LAPI handle.

### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hdl* passed in is not valid (not initialized or in terminated state).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Getcntr**, **LAPI\_Msgpoll**, **LAPI\_Nopoll\_wait**, **LAPI\_Waitcntr**

---

## LAPI\_Purge\_totask Subroutine

### Purpose

Allows a task to cancel messages to a given destination.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Purge_totask(hdl, dest)
```

```
lapi_handle_t hdl;
```

```
uint dest;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
int LAPI_PURGE_TOTASK(hndl, dest, ierror)
```

```
INTEGER hndl
```

```
INTEGER dest
```

```
INTEGER ierror
```

## Description

**Type of call:** recovery

This subroutine cancels messages and resets the state corresponding to messages in flight or submitted to be sent to a particular target task. This is an entirely local operation. For correct behavior a similar invocation is expected on the destination (if it exists). This function cleans up all the state associated with pending messages to the indicated target task. It is assumed that before the indicated task starts communicating with this task again, it also purges this instance (or that it was terminated and initialized again). It will also wake up all threads that are in **LAPI\_Nopoll\_wait** depending on how the arguments are passed to the **LAPI\_Nopoll\_wait** function. The behavior of **LAPI\_Purge\_totask** is undefined if LAPI collective functions are used.

**Note:** This subroutine should not be used when the parallel application is running in a PE/LoadLeveler environment.

**LAPI\_Purge\_totask** is normally used after connectivity has been lost between two tasks. If connectivity is restored, the tasks can be restored for LAPI communication by calling **LAPI\_Resume\_totask**.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*dest* Specifies the destination instance ID to which pending messages need to be cancelled.

### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

Use of this subroutine is *not* recommended on a system that is running Parallel Environment (PE).

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_TGT** Indicates that *dest* is outside the range of tasks defined in the job.

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Init**, **LAPI\_Nopoll\_wait**, **LAPI\_Resume\_totask**, **LAPI\_Term**

---

## LAPI\_Put Subroutine

### Purpose

Transfers data from a local task to a remote task.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Put(hndl, tgt, len, tgt_addr, org_addr, tgt_cntr, org_cntr, cmpl_cntr)
lapi_handle_t hndl;
uint          tgt;
ulong        len;
void         *tgt_addr;
void         *org_addr;
lapi_cntr_t  *tgt_cntr;
lapi_cntr_t  *org_cntr;
lapi_cntr_t  *cmpl_cntr;
```

### FORTRAN Syntax

```
include 'lapif.h'

int LAPI_PUT(hndl, tgt, len, tgt_addr, org_addr, tgt_cntr, org_cntr, ierror)
INTEGER hndl
INTEGER tgt
INTEGER (KIND=LAPI_LONG_TYPE) :: len
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_addr
INTEGER org_addr
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_cntr
TYPE (LAPI_CNTR_T) :: org_cntr
TYPE (LAPI_CNTR_T) :: cml_cntr
INTEGER ierror
```

### Description

**Type of call:** point-to-point communication (non-blocking)

Use this subroutine to transfer data from a local (origin) task to a remote (target) task. The origin counter will increment on the origin task upon origin buffer availability. The target counter will increment on the target and the completion counter will increment at the origin task upon message completion. Because there is no completion handler, message completion and target buffer availability are the same in this case.

This is a non-blocking call. The caller *cannot* assume that the data transfer has completed upon the return of the function. Instead, counters should be used to ensure correct buffer accesses as defined above.

Note that a zero-byte message does not transfer data, but it does have the same semantic with respect to counters as that of any other message.

### Parameters

#### INPUT

*hndl* Specifies the LAPI handle.



<i>tgt</i>	Specifies the task ID of the target task. The value of this parameter must be in the range $0 \leq tgt < \mathbf{NUM\_TASKS}$ .
<i>len</i>	Specifies the number of bytes to be transferred. This parameter must be in the range $0 \leq len \leq$ the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<i>tgt_addr</i>	Specifies the address on the target task where data is to be copied into. If <i>len</i> is <b>0</b> , The value of this parameter can be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<i>org_addr</i>	Specifies the address on the origin task from which data is to be copied. If <i>len</i> is <b>0</b> , The value of this parameter can be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).

## INPUT/OUTPUT

<i>tgt_cntr</i>	Specifies the target counter address. The target counter is incremented upon message completion. If this parameter is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), the target counter is not updated.
<i>org_cntr</i>	Specifies the origin counter address (in C) or the origin counter (in FORTRAN). The origin counter is incremented at buffer availability. If this parameter is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), the origin counter is not updated.
<i>compl_cntr</i>	Specifies the completion counter address (in C) or the completion counter (in FORTRAN) that is a reflection of <i>tgt_cntr</i> . The completion counter is incremented at the origin after <i>tgt_cntr</i> is incremented. If this parameter is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), the completion counter is not updated.

## OUTPUT

<i>ierror</i>	Specifies a FORTRAN return code. This is always the last parameter.
---------------	---

## C Examples

```
{
    /* initialize the table buffer for the data addresses          */
    /* get remote data buffer addresses                          */
    LAPI_Address_init(hndl, (void *)data_buffer, data_buffer_list);
    .
    .
    LAPI_Put(hndl, tgt, (ulong) data_len, (void *) (data_buffer_list[tgt]),
            (void *) data_buffer, tgt_cntr, org_cntr, compl_cntr);

    /* transfer data_len bytes from local address data buffer.   */
    /* write the data starting at address data_buffer_list[tgt] on */
    /* task tgt. tgt_cntr, org_cntr, and compl_cntr can be NULL. */
}
```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_DATA_LEN</b>	Indicates that the value of <i>len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_ORG_ADDR_NULL</b>	Indicates that the <i>org_addr</i> parameter passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but <i>len</i> is greater than <b>0</b> .

<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_ADDR_NULL</b>	Indicates that the <i>tgt_addr</i> parameter passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but <i>len</i> is greater than 0.
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Get**, **LAPI\_Getcptr**, **LAPI\_Qenv**, **LAPI\_Setcptr**, **LAPI\_Waitcptr**, **LAPI\_Xfer**

---

## LAPI\_Putv Subroutine

### Purpose

Transfers vectors of data from a local task to a remote task.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Putv(hdl, tgt, tgt_vec, org_vec, tgt_cptr, org_cptr, cmpl_cptr)
```

```
lapi_handle_t hdl;
uint tgt;
lapi_vec_t *tgt_vec;
lapi_vec_t *org_vec;
lapi_cptr_t *tgt_cptr;
lapi_cptr_t *org_cptr;
lapi_cptr_t *cmpl_cptr;
```

```
typedef struct {
    lapi_vectype_t vec_type; /* operation code */
    uint num_vecs; /* number of vectors */
    void **info; /* vector of information */
    ulong *len; /* vector of lengths */
} lapi_vec_t;
```

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_PUTV(hdl, tgt, tgt_vec, org_vec, tgt_cptr, org_cptr, cmpl_cptr, ierror)
```

```
INTEGER hdl
```

```
INTEGER tgt
```

```
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_vec
```

```
TYPE (LAPI_VEC_T) :: org_vec
```

```
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_cptr
```

```
TYPE (LAPI_CNTR_T) :: org_cptr
```

```
TYPE (LAPI_CNTR_T) :: cmpl_cptr
```

```
INTEGER ierror
```

The 32-bit version of the **LAPI\_VEC\_T** type is defined as:

```
TYPE LAPI_VEC_T
  SEQUENCE
  INTEGER(KIND = 4) :: vec_type
  INTEGER(KIND = 4) :: num_vecs
  INTEGER(KIND = 4) :: info
  INTEGER(KIND = 4) :: len
END TYPE LAPI_VEC_T
```

The 64-bit version of the **LAPI\_VEC\_T** type is defined as:

```
TYPE LAPI_VEC_T
  SEQUENCE
  INTEGER(KIND = 4) :: vec_type
  INTEGER(KIND = 4) :: num_vecs
  INTEGER(KIND = 8) :: info
  INTEGER(KIND = 8) :: len
END TYPE LAPI_VEC_T
```

## Description

**Type of call:** point-to-point communication (non-blocking)

**LAPI\_Putv** is the vector version of the **LAPI\_Put** call. Use this subroutine to transfer vectors of data from the origin task to the target task. The origin vector descriptions and the target vector descriptions are located in the address space of the *origin* task. However, the values specified in the *info* array of the target vector must be addresses in the address space of the *target* task.

The calling program *cannot* assume that the origin buffer can be changed or that the contents of the target buffers on the target task are ready for use upon function return. After the origin counter (*org\_cntr*) is incremented, the origin buffers can be modified by the origin task. After the target counter (*tgt\_cntr*) is incremented, the target buffers can be modified by the target task. If you provide a completion counter (*cmp\_cntr*), it is incremented at the origin after the target counter (*tgt\_cntr*) has been incremented at the target. If the values of any of the counters or counter addresses are NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the data transfer occurs, but the corresponding counter increments do not occur.

If a strided vector is being transferred, the size of each block must not be greater than the stride size in bytes.

The length of any vector pointed to by *org\_vec* must be equal to the length of the corresponding vector pointed to by *tgt\_vec*.

LAPI does not check for any overlapping regions among vectors either at the origin or the target. If the overlapping regions exist on the target side, the contents of the target buffer are undefined after the operation.

See **LAPI\_Amsendv** for more information about using the various vector types. (**LAPI\_Putv** does not support the **LAPI\_GEN\_GENERIC** type.)

## Parameters

### INPUT

<i>hndl</i>	Specifies the LAPI handle.
<i>tgt</i>	Specifies the task ID of the target task. The value of this parameter must be in the range 0 <= <i>tgt</i> < <b>NUM_TASKS</b> .
<i>tgt_vec</i>	Points to the target vector description.
<i>org_vec</i>	Points to the origin vector description.

## INPUT/OUTPUT

- tgt\_cntr* Specifies the target counter address. The target counter is incremented upon message completion. If this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the target counter is not updated.
- org\_cntr* Specifies the origin counter address (in C) or the origin counter (in FORTRAN). The origin counter is incremented at buffer availability. If this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the origin counter is not updated.
- compl\_cntr* Specifies the completion counter address (in C) or the completion counter (in FORTRAN) that is a reflection of *tgt\_cntr*. The completion counter is incremented at the origin after *tgt\_cntr* is incremented. If this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), the completion counter is not updated.

## OUTPUT

- ierror* Specifies a FORTRAN return code. This is always the last parameter.

## C Examples

To put a LAPI\_GEN\_IOVECTOR:

```
{
    /* retrieve a remote data buffer address for data to transfer, */
    /* such as through LAPI_Address_init */

    /* task that calls LAPI_Putv sets up both org_vec and tgt_vec */
    org_vec->num_vecs = NUM_VECS;
    org_vec->vec_type = LAPI_GEN_IOVECTOR;
    org_vec->len      = (unsigned long *)
    malloc(NUM_VECS*sizeof(unsigned long));
    org_vec->info    = (void **) malloc(NUM_VECS*sizeof(void *));

    /* each org_vec->info[i] gets a base address on the origin task */
    /* each org_vec->len[i] gets the number of bytes to transfer */
    /* from org_vec->info[i] */

    tgt_vec->num_vecs = NUM_VECS;
    tgt_vec->vec_type = LAPI_GEN_IOVECTOR;
    tgt_vec->len      = (unsigned long *)
    malloc(NUM_VECS*sizeof(unsigned long));
    tgt_vec->info    = (void **) malloc(NUM_VECS*sizeof(void *));

    /* each tgt_vec->info[i] gets a base address on the target task */
    /* each tgt_vec->len[i] gets the number of bytes to write to vec->info[i] */
    /* For LAPI_GEN_IOVECTOR, num_vecs, vec_type, and len must be the same */

    LAPI_Putv(hndl, tgt, tgt_vec, org_vec, tgt_cntr, org_cntr, compl_cntr);
    /* tgt_cntr, org_cntr and compl_cntr can all be NULL */

    /* data will be transferred as follows: */
    /* org_vec->len[0] bytes will be retrieved from */
    /* org_vec->info[0] and written to tgt_vec->info[0] */
    /* org_vec->len[1] bytes will be retrieved from */
    /* org_vec->info[1] and written to tgt_vec->info[1] */
    .
    .
    .
    /* org_vec->len[NUM_VECS-1] bytes will be retrieved */
    /* from org_vec->info[NUM_VECS-1] and written to */
    /* tgt_vec->info[NUM_VECS-1] */
}

```

See the example in **LAPI\_Amsendv** for information on other vector types.

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_ORG_EXTENT</b>	Indicates that the <i>org_vec</i> 's extent (stride * <i>num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_ORG_STRIDE</b>	Indicates that the <i>org_vec</i> stride is less than block.
<b>LAPI_ERR_ORG_VEC_ADDR</b>	Indicates that the <i>org_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but its length ( <i>org_vec-&gt;len[i]</i> ) is not 0.
<b>LAPI_ERR_ORG_VEC_LEN</b>	Indicates that the sum of <i>org_vec-&gt;len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_ORG_VEC_NULL</b>	Indicates that the <i>org_vec</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_ORG_VEC_TYPE</b>	Indicates that the <i>org_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_STRIDE_ORG_VEC_ADDR_NULL</b>	Indicates that the strided vector address <i>org_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_STRIDE_TGT_VEC_ADDR_NULL</b>	Indicates that the strided vector address <i>tgt_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_EXTENT</b>	Indicates that <i>tgt_vec</i> 's extent (stride * <i>num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.
<b>LAPI_ERR_TGT_STRIDE</b>	Indicates that the <i>tgt_vec</i> stride is less than block.
<b>LAPI_ERR_TGT_VEC_ADDR</b>	Indicates that the <i>tgt_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but its length ( <i>tgt_vec-&gt;len[i]</i> ) is not 0.
<b>LAPI_ERR_TGT_VEC_LEN</b>	Indicates that the sum of <i>tgt_vec-&gt;len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_TGT_VEC_NULL</b>	Indicates that <i>tgt_vec</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT_VEC_TYPE</b>	Indicates that the <i>tgt_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_VEC_LEN_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different lengths ( <i>len[ ]</i> ).
<b>LAPI_ERR_VEC_NUM_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different <i>num_vecs</i> .
<b>LAPI_ERR_VEC_TYPE_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different vector types ( <i>vec_type</i> ).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Amsendv**, **LAPI\_Getcntr**, **LAPI\_Getv**, **LAPI\_Qenv**, **LAPI\_Setcntr**, **LAPI\_Waitcntr**, **LAPI\_Xfer**

---

## LAPI\_Qenv Subroutine

### Purpose

Used to query LAPI for runtime task information.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapif.h>
```

```
int LAPI_Qenv(hndl, query, ret_val)
lapi_handle_t hndl;
lapi_query_t query;
int *ret_val; /* ret_val's type varies (see Additional query types) */
```

### FORTTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_QENV(hndl, query, ret_val, ierror)
INTEGER hndl
INTEGER query
INTEGER ret_val /* ret_val's type varies (see Additional query types) */
INTEGER ierror
```

### Description

**Type of call:** local queries

Use this subroutine to query runtime settings and statistics from LAPI. LAPI defines a set of query types as an enumeration in **lapi.h** for C and explicitly in the 32-bit and 64-bit versions of **lapif.h** for FORTRAN.

For example, you can query the size of the table that LAPI uses for the **LAPI\_Addr\_set** subroutine using a *query* value of **LOC\_ADDRTBL\_SZ**:

```
LAPI_Qenv(hndl, LOC_ADDRTBL_SZ, &ret_val);
```

*ret\_val* will contain the upper bound on the table index. A subsequent call to **LAPI\_Addr\_set** (*hndl*, *addr*, *addr\_hndl*); could then ensure that the value of *addr\_hndl* is between **0** and *ret\_val*.

When used to show the size of a parameter, a comparison of values, or a range of values, valid values for the *query* parameter of the **LAPI\_Qenv** subroutine appear in **SMALL, BOLD** capital letters. For example:

### NUM\_TASKS

is a shorthand notation for:

```
LAPI_Qenv(hndl, NUM_TASKS, ret_val)
```

In C, `lapi_query_t` defines the valid types of LAPI queries:

```
typedef enum {
    TASK_ID=0,      /* Query the task ID of the current task in the job */
    NUM_TASKS,     /* Query the number of tasks in the job */
    MAX_UHDR_SZ,   /* Query the maximum user header size for active messaging */
    MAX_DATA_SZ,  /* Query the maximum data length that can be sent */
    ERROR_CHK,    /* Query and set parameter checking on (1) or off (0) */
    TIMEOUT,      /* Query and set the current communication timeout setting */
                 /* in seconds */
    MIN_TIMEOUT, /* Query the minimum communication timeout setting in seconds */
    MAX_TIMEOUT, /* Query the maximum communication timeout setting in seconds */
    INTERRUPT_SET, /* Query and set interrupt mode on (1) or off (0) */
    MAX_PORTS,    /* Query the maximum number of available communication ports */
    MAX_PKT_SZ,  /* This is the payload size of 1 packet */
    NUM_REX_BUFS, /* Number of retransmission buffers */
    REX_BUF_SZ,  /* Size of each retransmission buffer in bytes */
    LOC_ADDR_TBL_SZ, /* Size of address store table used by LAPI_Addr_set */
    EPOCH_NUM,   /* No longer used by LAPI (supports legacy code) */
    USE_THRESH,  /* No longer used by LAPI (supports legacy code) */
    RCV_FIFO_SIZE, /* No longer used by LAPI (supports legacy code) */
    MAX_ATOM_SIZE, /* Query the maximum atom size for a DGSP accumulate transfer */
    BUF_CP_SIZE, /* Query the size of the message buffer to save (default 128b) */
    MAX_PKTS_OUT, /* Query the maximum number of messages outstanding /
                 /* destination */
    ACK_THRESHOLD, /* Query and set the threshold of acknowledgments going
                 /* back to the source */
    QUERY_SHM_ENABLED, /* Query to see if shared memory is enabled */
    QUERY_SHM_NUM_TASKS, /* Query to get the number of tasks that use shared
                 /* memory */
    QUERY_SHM_TASKS, /* Query to get the list of task IDs that make up shared
                 /* memory; pass in an array of size QUERY_SHM_NUM_TASKS */
    QUERY_STATISTICS, /* Query to get packet statistics from LAPI, as
                 /* defined by the lapi_statistics_t structure. For
                 /* this query, pass in 'lapi_statistics_t *' rather
                 /* than 'int *ret_val'; otherwise, the data will
                 /* overflow the buffer. */
    PRINT_STATISTICS, /* Query debug print function to print out statistics */
    QUERY_SHM_STATISTICS, /* Similar query as QUERY_STATISTICS for shared
                 /* memory path. */
    QUERY_LOCAL_SEND_STATISTICS, /* Similar query as QUERY_STATISTICS
                 /* for local copy path. */
    BULK_XFER, /* Query to see if bulk transfer is enabled (1) or disabled (0) */
    BULK_MIN_MSG_SIZE, /* Query the current bulk transfer minimum message size */
    LAST_QUERY
} lapi_query_t;
```

```
typedef struct {
    lapi_long_t Tot_dup_pkt_cnt; /* Total duplicate packet count */
    lapi_long_t Tot_retrans_pkt_cnt; /* Total retransmit packet count */
    lapi_long_t Tot_gho_pkt_cnt; /* Total Ghost packet count */
    lapi_long_t Tot_pkt_sent_cnt; /* Total packet sent count */
    lapi_long_t Tot_pkt_rcv_cnt; /* Total packet receive count */
    lapi_long_t Tot_data_sent; /* Total data sent */
    lapi_long_t Tot_data_rcv; /* Total data receive */
} lapi_statistics_t;
```

In FORTRAN, the valid types of LAPI queries are defined in `lapif.h` as follows:

```
integer TASK_ID,NUM_TASKS,MAX_UHDR_SZ,MAX_DATA_SZ,ERROR_CHK
integer TIMEOUT,MIN_TIMEOUT,MAX_TIMEOUT
integer INTERRUPT_SET,MAX_PORTS,MAX_PKT_SZ,NUM_REX_BUFS
integer REX_BUF_SZ,LOC_ADDR_TBL_SZ,EPOCH_NUM,USE_THRESH
integer RCV_FIFO_SIZE,MAX_ATOM_SIZE,BUF_CP_SIZE
integer MAX_PKTS_OUT,ACK_THRESHOLD,QUERY_SHM_ENABLED
```

```

integer QUERY_SHM_NUM_TASKS,QUERY_SHM_TASKS
integer QUERY_STATISTICS,PRINT_STATISTICS
integer QUERY_SHM_STATISTICS,QUERY_LOCAL_SEND_STATISTICS
integer BULK_XFER,BULK_MIN_MSG_SIZE,
integer LAST_QUERY
parameter (TASK_ID=0,NUM_TASKS=1,MAX_UHDR_SZ=2,MAX_DATA_SZ=3)
parameter (ERROR_CHK=4,TIMEOUT=5,MIN_TIMEOUT=6)
parameter (MAX_TIMEOUT=7,INTERRUPT_SET=8,MAX_PORTS=9)
parameter (MAX_PKT_SZ=10,NUM_REX_BUFS=11,REX_BUF_SZ=12)
parameter (LOC_ADDRTBL_SZ=13,EPOCH_NUM=14,USE_THRESH=15)
parameter (RCV_FIFO_SIZE=16,MAX_ATOM_SIZE=17,BUF_CP_SIZE=18)
parameter (MAX_PKTS_OUT=19,ACK_THRESHOLD=20)
parameter (QUERY_SHM_ENABLED=21,QUERY_SHM_NUM_TASKS=22)
parameter (QUERY_SHM_TASKS=23,QUERY_STATISTICS=24)
parameter (PRINT_STATISTICS=25)
parameter (QUERY_SHM_STATISTICS=26,QUERY_LOCAL_SEND_STATISTICS=27)
parameter (BULK_XFER=28,BULK_MIN_MSG_SIZE=29)
parameter (LAST_QUERY=30)

```

## Additional query types

LAPI provides additional query types for which the behavior of **LAPI\_Qenv** is slightly different:

### PRINT\_STATISTICS

When passed this query type, LAPI sends data transfer statistics to standard output. In this case, *ret\_val* is unaffected. However, LAPI's error checking requires that the value of *ret\_val* is not NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN) for all **LAPI\_Qenv** types (including **PRINT\_STATISTICS**).

### QUERY\_LOCAL\_SEND\_STATISTICS

When passed this query type, **LAPI\_Qenv** interprets *ret\_val* as a pointer to type **lapi\_statistics\_t**. Upon function return, the fields of the structure contain LAPI's data transfer statistics for data transferred through intra-task local copy. The packet count will be 0.

### QUERY\_SHM\_STATISTICS

When passed this query type, **LAPI\_Qenv** interprets *ret\_val* as a pointer to type **lapi\_statistics\_t**. Upon function return, the fields of the structure contain LAPI's data transfer statistics for data transferred through shared memory.

### QUERY\_SHM\_TASKS

When passed this query type, **LAPI\_Qenv** returns a list of task IDs with which this task can communicate using shared memory. *ret\_val* must be an **int \*** with enough space to hold **NUM\_TASKS** integers. For each task *i*, if it is possible to use shared memory, *ret\_val[i]* will contain the shared memory task ID. If it is not possible to use shared memory, *ret\_val[i]* will contain **-1**.

### QUERY\_STATISTICS

When passed this query type, **LAPI\_Qenv** interprets *ret\_val* as a pointer to type **lapi\_statistics\_t**. Upon function return, the fields of the structure contain LAPI's data transfer statistics for data transferred using the user space (US) protocol or UDP/IP.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*query* Specifies the type of query you want to request. In C, the values for *query* are defined by the **lapi\_query\_t** enumeration in **lapi.h**. In FORTRAN, these values are defined explicitly in the 32-bit version and the 64-bit version of **lapif.h**.

### OUTPUT



*ret\_val* Specifies the reference parameter for LAPI to store as the result of the query. The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Return values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_QUERY\_TYPE** Indicates that the query passed in is not valid.

**LAPI\_ERR\_RET\_PTR\_NULL** Indicates that the value of the *ret\_val* pointer is NULL (in C) or that the value of *ret\_val* is **LAPI\_ADDR\_NULL** (in FORTRAN).

## C Examples

To query runtime values from LAPI:

```
{
    int          task_id;
    lapi_statistics_t stats;
    .
    .
    .
    LAPI_Qenv(hndl, TASK_ID, &task_id);
    /* task_id now contains the task ID */
    .
    .
    .
    LAPI_Qenv(hndl, QUERY_STATISTICS, (int *)&stats);
    /* the fields of the stats structure are now
       filled in with runtime values */
    .
    .
    .
}
```

## Location

**/usr/lib/liblapi\_r.a**

## Related Information

Subroutines: **LAPI\_Amsend**, **LAPI\_Get**, **LAPI\_Put**, **LAPI\_Senv**, **LAPI\_Xfer**

---

## LAPI\_Resume\_totask Subroutine

### Purpose

Re-enables the sending of messages to the task.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Resume_totask(hndl, dest)
lapi_handle_t hndl;
uint dest;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
int LAPI_RESUME_TOTASK(hndl, dest, ierror)  
INTEGER hndl  
INTEGER dest  
INTEGER ierror
```

## Description

**Type of call:** recovery

This subroutine is used in conjunction with **LAPI\_Purge\_totask**. It enables LAPI communication to be reestablished for a task that had previously been purged. The purged task must either restart LAPI or execute a **LAPI\_Purge\_totask/LAPI\_Resume\_totask** sequence for this task.

## Parameters

### INPUT

*hndl*                Specifies the LAPI handle.  
*dest*                Specifies the destination instance ID with which to resume communication.

### OUTPUT

*ierror*              Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

Use of this subroutine is *not* recommended on a system that is running Parallel Environment (PE).

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Init**, **LAPI\_Nopoll\_wait**, **LAPI\_Purge\_totask**, **LAPI\_Term**

---

## LAPI\_Rmw Subroutine

### Purpose

Provides data synchronization primitives.

### Library

Availability Library (**liblapi\_r.a**)

## C Syntax

```
#include <lapi.h>
```

```
int LAPI_Rmw(hdl, op, tgt, tgt_var, in_val, prev_tgt_val, org_cntr)
```

```
lapi_handle_t hdl;  
RMW_ops_t op;  
uint tgt;  
int *tgt_var;  
int *in_val;  
int *prev_tgt_val;  
lapi_cntr_t *org_cntr;
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_RMW(hdl, op, tgt, tgt_var, in_val, prev_tgt_val, org_cntr, ierror)  
INTEGER hdl  
INTEGER op  
INTEGER tgt  
INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_var  
INTEGER in_val  
INTEGER prev_tgt_val  
TYPE (LAPI_CNTR_T) :: org_cntr  
INTEGER ierror
```

## Description

**Type of call:** point-to-point communication (non-blocking)

Use this subroutine to synchronize two independent pieces of data, such as two tasks sharing a common data structure. The operation is performed at the target task (*tgt*) and is atomic. The operation takes an input value (*in\_val*) from the origin and performs one of four operations (*op*) on a variable (*tgt\_var*) at the target (*tgt*), and then replaces the target variable (*tgt\_var*) with the results of the operation (*op*). The original value (*prev\_tgt\_val*) of the target variable (*tgt\_var*) is returned to the origin.

The operations (*op*) are performed over the context referred to by *hdl*. The outcome of the execution of these calls is as if the following code was executed atomically:

```
*prev_tgt_val = *tgt_var;  
*tgt_var      = f(*tgt_var, *in_val);
```

where:

f(a,b) = a + b for **FETCH\_AND\_ADD**

f(a,b) = a | b for **FETCH\_AND\_OR** (bitwise or)

f(a,b) = b for **SWAP**

For **COMPARE\_AND\_SWAP**, *in\_val* is treated as a pointer to an array of two integers, and the *op* is the following atomic operation:

```
if(*tgt_var == in_val[0]) {  
    *prev_tgt_val = TRUE;  
    *tgt_var      = in_val[1];  
} else {  
    *prev_tgt_val = FALSE;  
}
```

All **LAPI\_Rmw** calls are non-blocking. To test for completion, use the **LAPI\_Getcntr** and **LAPI\_Waitcntr** subroutines. **LAPI\_Rmw** does not include a target counter (*tgt\_cntr*), so **LAPI\_Rmw** calls do not provide any indication of completion on the target task (*tgt*).

## Parameters

### INPUT

<i>hndl</i>	Specifies the LAPI handle.
<i>op</i>	Specifies the operation to be performed. The valid operations are: <ul style="list-style-type: none"> <li>• <b>COMPARE_AND_SWAP</b></li> <li>• <b>FETCH_AND_ADD</b></li> <li>• <b>FETCH_AND_OR</b></li> <li>• <b>SWAP</b></li> </ul>
<i>tgt</i>	Specifies the task ID of the target task where the read-modify-write (Rmw) variable resides. The value of this parameter must be in the range $0 \leq \textit{tgt} < \textit{NUM\_TASKS}$ .
<i>tgt_var</i>	Specifies the target read-modify-write (Rmw) variable (in FORTRAN) or its address (in C). The value of this parameter cannot be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<i>in_val</i>	Specifies the value that is passed in to the operation ( <i>op</i> ). This value cannot be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).

### INPUT/OUTPUT

<i>prev_tgt_val</i>	Specifies the location at the origin in which the previous <i>tgt_var</i> on the target task is stored before the operation ( <i>op</i> ) is executed. The value of this parameter can be NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<i>org_cntr</i>	Specifies the origin counter address (in C) or the origin counter (in FORTRAN). If <i>prev_tgt_val</i> is set, the origin counter ( <i>org_cntr</i> ) is incremented when <i>prev_tgt_val</i> is returned to the origin side. If <i>prev_tgt_val</i> is not set, the origin counter ( <i>org_cntr</i> ) is updated after the operation ( <i>op</i> ) is completed at the target side.

### OUTPUT

<i>ierror</i>	Specifies a FORTRAN return code. This is always the last parameter.
---------------	---

## Restrictions

LAPI statistics are *not* reported for shared memory communication and data transfer, or for messages that a task sends to itself.

## C Examples

1. To synchronize a data value between two tasks (with **FETCH\_AND\_ADD**):

```
{
    int local_var;
    int *addr_list;

    /* both tasks initialize local_var to a value */
    /* local_var addresses are exchanged and stored */
    /* in addr_list (using LAPI_Address_init). */
    /* addr_list[tgt] now contains the address of */
    /* local_var on tgt */
    .
    .
    .
    /* add value to local_var on some task */
}
```

```

    /* use LAPI to add value to local_var on remote task */
    LAPI_Rmw(hndl, FETCH_AND_ADD, tgt, addr_list[tgt],
             value, prev_tgt_val, &org_cntr);

    /* local_var on the remote task has been increased */
    /* by value. prev_tgt_val now contains the value */
    /* of local_var on remote task before the addition */
}

```

2. To synchronize a data value between two tasks (with **SWAP**):

```

{
    int local_var;
    int *addr_list;

    /* local_var addresses are exchanged and stored */
    /* in addr_list (using LAPI_Address_init). */
    /* addr_list[tgt] now contains the address of */
    /* local_var on tgt. */
    .
    .
    .
    /* local_var is assigned some value */

    /* assign local_var to local_var on remote task */
    LAPI_Rmw(hndl, SWAP, tgt, addr_list[tgt],
             local_var, prev_tgt_val, &org_cntr);

    /* local_var on the remote task is now equal to */
    /* local_var on the local task. prev_tgt_val now */
    /* contains the value of local_var on the remote */
    /* task before the swap. */
}

```

3. To conditionally swap a data value (with **COMPARE\_AND\_SWAP**):

```

{
    int local_var;
    int *addr_list;
    int in_val[2];

    /* local_var addresses are exchanged and stored */
    /* in addr_list (using LAPI_Address_init). */
    /* addr_list[tgt] now contains the address of */
    /* local_var on tgt. */
    .
    .
    .
    /* if local_var on remote_task is equal to comparator, */
    /* assign value to local_var on remote task */

    in_val[0] = comparator;
    in_val[1] = value;

    LAPI_Rmw(hndl, COMPARE_AND_SWAP, tgt, addr_list[tgt],
             in_val, prev_tgt_val, &org_cntr);

    /* local_var on the remote task is now in_val[1] if it */
    /* had previously been equal to in_val[0]. If the swap */
    /* was performed, prev_tgt_val now contains TRUE; */
    /* otherwise, it contains FALSE. */
}

```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_IN_VAL_NULL</b>	Indicates that the <i>in_val</i> pointer is NULL (in C) or that the value of <i>in_val</i> is <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_RMW_OP</b>	Indicates that <i>op</i> is not valid.
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.
<b>LAPI_ERR_TGT_VAR_NULL</b>	Indicates that the <b>tgt_var</b> address is NULL (in C) or that the value of <b>tgt_var</b> is <b>LAPI_ADDR_NULL</b> (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address\_init**, **LAPI\_Getcctr**, **LAPI\_Qenv**, **LAPI\_Rmw64**, **LAPI\_Setcctr**, **LAPI\_Waitcctr**, **LAPI\_Xfer**

---

## LAPI\_Rmw64 Subroutine

### Purpose

Provides data synchronization primitives for 64-bit applications.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Rmw64(hndl, op, tgt, tgt_var, in_val, prev_tgt_val, org_cctr)
```

```
lapi_handle_t hndl;  
Rmw_ops_t op;  
uint tgt;  
long long *tgt_var;  
long long *in_val;  
long long *prev_tgt_val;  
lapi_cctr_t *org_cctr;
```

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_RMW64(hndl, op, tgt, tgt_var, in_val, prev_tgt_val, org_cctr, ierror)
```

```
INTEGER hndl  
INTEGER op  
INTEGER tgt
```

```

INTEGER (KIND=LAPI_ADDR_TYPE) :: tgt_var
INTEGER (KIND=LAPI_LONG_LONG_TYPE) :: in_val, prev_tgt_val
TYPE (LAPI_CNTR_T) :: org_cntr
INTEGER ierror

```

## Description

**Type of call:** point-to-point communication (non-blocking)

This subroutine is the 64-bit version of **LAPI\_Rmw**. It is used to synchronize two independent pieces of 64-bit data, such as two tasks sharing a common data structure. The operation is performed at the target task (*tgt*) and is atomic. The operation takes an input value (*in\_val*) from the origin and performs one of four operations (*op*) on a variable (*tgt\_var*) at the target (*tgt*), and then replaces the target variable (*tgt\_var*) with the results of the operation (*op*). The original value (*prev\_tgt\_val*) of the target variable (*tgt\_var*) is returned to the origin.

The operations (*op*) are performed over the context referred to by *hndl*. The outcome of the execution of these calls is as if the following code was executed atomically:

```

*prev_tgt_val = *tgt_var;
*tgt_var      = f(*tgt_var, *in_val);

```

where:

f(a,b) = a + b for **FETCH\_AND\_ADD**

f(a,b) = a | b for **FETCH\_AND\_OR** (bitwise or)

f(a,b) = b for **SWAP**

For **COMPARE\_AND\_SWAP**, *in\_val* is treated as a pointer to an array of two integers, and the *op* is the following atomic operation:

```

if(*tgt_var == in_val[0]) {
    *prev_tgt_val = TRUE;
    *tgt_var      = in_val[1];
} else {
    *prev_tgt_val = FALSE;
}

```

This subroutine can also be used on a 32-bit processor.

All **LAPI\_Rmw64** calls are non-blocking. To test for completion, use the **LAPI\_Getcntr** and **LAPI\_Waitcntr** subroutines. **LAPI\_Rmw64** does not include a target counter (*tgt\_cntr*), so **LAPI\_Rmw64** calls do not provide any indication of completion on the target task (*tgt*).

## Parameters

### INPUT

- |             |  |
|-------------|--|
| <i>hndl</i> | Specifies the LAPI handle.   |
| <i>op</i>   | Specifies the operation to be performed. The valid operations are: <ul style="list-style-type: none"> <li>• <b>COMPARE_AND_SWAP</b></li> <li>• <b>FETCH_AND_ADD</b></li> <li>• <b>FETCH_AND_OR</b></li> <li>• <b>SWAP</b></li> </ul> |

*tgt* Specifies the task ID of the target task where the read-modify-write (Rmw64) variable resides. The value of this parameter must be in the range  $0 \leq tgt < \mathbf{NUM\_TASKS}$ .

*tgt\_var* Specifies the target read-modify-write (Rmw64) variable (in FORTRAN) or its address (in C). The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*in\_val* Specifies the value that is passed in to the operation (*op*). This value cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

## INPUT/OUTPUT

*prev\_tgt\_val* Specifies the location at the origin in which the previous *tgt\_var* on the target task is stored before the operation (*op*) is executed. The value of this parameter can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*org\_cntr* Specifies the origin counter address (in C) or the origin counter (in FORTRAN). If *prev\_tgt\_val* is set, the origin counter (*org\_cntr*) is incremented when *prev\_tgt\_val* is returned to the origin side. If *prev\_tgt\_val* is not set, the origin counter (*org\_cntr*) is updated after the operation (*op*) is completed at the target side.

## OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

LAPI statistics are *not* reported for shared memory communication and data transfer, or for messages that a task sends to itself.

## C Examples

1. To synchronize a data value between two tasks (with **FETCH\_AND\_ADD**):

```
{
    long long local_var;
    long long *addr_list;

    /* both tasks initialize local_var to a value */
    /* local_var addresses are exchanged and stored */
    /* in addr_list (using LAPI_Address_init64) */
    /* addr_list[tgt] now contains address of */
    /* local_var on tgt */
    .
    .
    .
    /* add value to local_var on some task */

    /* use LAPI to add value to local_var on remote task */
    LAPI_Rmw64(hnd1, FETCH_AND_ADD, tgt, addr_list[tgt],
               value, prev_tgt_val, &org_cntr);

    /* local_var on remote task has been increased */
    /* by value. prev_tgt_val now contains value of */
    /* local_var on remote task before the addition */
}
```

2. To synchronize a data value between two tasks (with **SWAP**):

```
{
    long long local_var;
    long long *addr_list;
```



```

/* local_var addresses are exchanged and stored      */
/* in addr_list (using LAPI_Address_init64).        */
/* addr_list[tgt] now contains the address of      */
/* local_var on tgt.                               */
.
.
.
/* local_var is assigned some value                */

/* assign local_var to local_var on the remote task */
LAPI_Rmw64(hndl, SWAP, tgt, addr_list[tgt],
           local_var, prev_tgt_val, &org_cntr);

/* local_var on the remote task is now equal to local_var */
/* on the local task. prev_tgt_val now contains the value */
/* of local_var on the remote task before the swap.      */
}

```

### 3. To conditionally swap a data value (with **COMPARE\_AND\_SWAP**):

```

{
    long long local_var;
    long long *addr_list;
    long long in_val[2];

    /* local_var addresses are exchanged and stored      */
    /* in addr_list (using LAPI_Address_init64).        */
    /* addr_list[tgt] now contains the address of      */
    /* local_var on tgt.                               */
    .
    .
    .
    /* if local_var on remote_task is equal to comparator, */
    /* assign value to local_var on the remote task        */
    in_val[0] = comparator;
    in_val[1] = value;

    LAPI_Rmw64(hndl, COMPARE_AND_SWAP, tgt, addr_list[tgt],
               in_val, prev_tgt_val, &org_cntr);

    /* local_var on remote task is now in_val[1] if it    */
    /* had previously been equal to in_val[0]. If the    */
    /* swap was performed, prev_tgt_val now contains     */
    /* TRUE; otherwise, it contains FALSE.              */
}

```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_IN_VAL_NULL</b>	Indicates that the <i>in_val</i> pointer is NULL (in C) or that the value of <i>in_val</i> is <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_RMW_OP</b>	Indicates that <i>op</i> is not valid.
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.

**LAPI\_ERR\_TGT\_VAR\_NULL** Indicates that the **tgt\_var** address is NULL (in C) or that the value of **tgt\_var** is **LAPI\_ADDR\_NULL** (in FORTRAN).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Address\_init64**, **LAPI\_Getcptr**, **LAPI\_Qenv**, **LAPI\_Rmw**, **LAPI\_Setcptr**, **LAPI\_Waitcptr**, **LAPI\_Xfer**

---

## LAPI\_Senv Subroutine

### Purpose

Used to set a runtime variable.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Senv(hdl, query, set_val)
lapi_handle_t hdl;
lapi_query_t query;
int set_val;
```

### FORTRAN Syntax

```
include 'lapi.h'
```

```
LAPI_SENV(hdl, query, set_val, ierror)
INTEGER hdl
INTEGER query
INTEGER set_val
INTEGER ierror
```

### Description

**Type of call:** local queries

Use this subroutine to set runtime attributes for a specific LAPI instance. In C, the **lapi\_query\_t** enumeration defines the attributes that can be set at runtime. These attributes are defined explicitly in FORTRAN. See **LAPI\_Qenv** for more information.

You can use **LAPI\_Senv** to set these runtime attributes: **ACK\_THRESHOLD**, **ERROR\_CHK**, **INTERRUPT\_SET**, and **TIMEOUT**.

### Parameters

#### INPUT

*hdl* Specifies the LAPI handle.

*query* Specifies the type of query that you want to set. In C, the values for *query* are defined by the **lapi\_query\_t** enumeration in **lapi.h**. In FORTRAN, these values are defined explicitly in the 32-bit version and the 64-bit version of **lapi.h**.

*set\_val* Specifies the integer value of the query that you want to set.

## OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

LAPI statistics are *not* reported for shared memory communication and data transfer, or for messages that a task sends to itself.

## C Examples

The following values can be set using **LAPI\_Senv**:

```
ACK_THRESHOLD:
int value;
LAPI_Senv(hndl, ACK_THRESHOLD, value);
/* LAPI sends packet acknowledgements (acks) in groups, waiting until */
/* ACK_THRESHOLD packets have arrived before returning a group of acks */
/* The valid range for ACK_THRESHOLD is (1 <= value <= 30) */
/* The default is 30. */

ERROR_CHK:
boolean toggle;
LAPI_Senv(hndl, ERROR_CHK, toggle);
/* Indicates whether LAPI should perform error checking. If set, LAPI */
/* calls will perform bounds-checking on parameters. Error checking */
/* is disabled by default. */

INTERRUPT_SET:
boolean toggle;
LAPI_Senv(hndl, INTERRUPT_SET, toggle);
/* Determines whether LAPI will respond to interrupts. If interrupts */
/* are disabled, LAPI will poll for message completion. */
/* toggle==True will enable interrupts, False will disable. */
/* Interrupts are enabled by default. */

TIMEOUT:
int value;
LAPI_Senv(hndl, TIMEOUT, value);
/* LAPI will time out on a communication if no response is received */
/* within timeout seconds. Valid range is (10 <= timeout <= 86400). */
/* 86400 seconds = 24 hours. Default value is 900 (15 minutes). */
```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hndl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_QUERY_TYPE</b>	Indicates the query passed in is not valid.
<b>LAPI_ERR_SET_VAL</b>	Indicates the <i>set_val</i> pointer is not in valid range.

## Location

`/usr/lib/liblapi_r.a`

## Related Information

Subroutines: **LAPI\_Qenv**

---

## LAPI\_Setcntr Subroutine

### Purpose

Used to set a counter to a specified value.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Setcntr(hdl, cntr, val)
lapi_handle_t hdl;
lapi_cntr_t *cntr;
int val;
```

### FORTRAN Syntax

```
include 'lapif.h'

LAPI_SETCNTR(hdl, cntr, val, ierror)
INTEGER hdl
TYPE (LAPI_CNTR_T) :: cntr
INTEGER val
INTEGER ierror
```

### Description

**Type of call:** Local counter manipulation

This subroutine sets *cntr* to the value specified by *val*. Because the **LAPI\_Getcntr/LAPI\_Setcntr** sequence cannot be made atomic, you should only use **LAPI\_Setcntr** when you know there will not be any competing operations.

### Parameters

#### INPUT

*hdl* Specifies the LAPI handle.  
*val* Specifies the value to which the counter needs to be set.

#### INPUT/OUTPUT

*cntr* Specifies the address of the counter to be set (in C) or the counter structure (in FORTRAN). The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

#### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

### Restrictions

LAPI statistics are *not* reported for shared memory communication and data transfer, or for messages that a task sends to itself.

### C Examples

To initialize a counter for use in a communication API call:

```

{
  lapi_cntr_t  my_tgt_cntr, *tgt_cntr_array;
  int         initial_value, expected_value, current_value;
  lapi_handle_t hndl;
  .
  .
  .
  /*
   * Note: the code below is executed on all tasks
   */

  /* initialize, allocate and create structures */
  initial_value = 0;
  expected_value = 1;

  /* set the cntr to zero */
  LAPI_Setcntr(hndl, &my_tgt_cntr, initial_value);
  /* set other counters */
  .
  .
  .
  /* exchange counter addresses, LAPI_Address_init synchronizes */
  LAPI_Address_init(hndl, &my_tgt_cntr, tgt_cntr_array);
  /* more address exchanges */
  .
  .
  .
  /* Communication calls using my_tgt_cntr */
  LAPI_Put(....., tgt_cntr_array[tgt], ....);
  .
  .
  .
  /* Wait for counter to reach value */
  for (;;) {
    LAPI_Getcntr(hndl, &my_tgt_cntr, &current_value);
    if (current_value >= expected_value) {
      break; /* out of infinite loop */
    } else {
      LAPI_Probe(hndl);
    }
  }
  .
  .
  .
  /* Quiesce/synchronize to ensure communication using our counter is done */
  LAPI_Gfence(hndl);
  /* Reset the counter */
  LAPI_Setcntr(hndl, &my_tgt_cntr, initial_value);
  /*
   * Synchronize again so that no other communication using the counter can
   * begin from any other task until we're all finished resetting the counter.
   */
  LAPI_Gfence(hndl);

  /* More communication calls */
  .
  .
  .
}

```

## Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_CNTR_NULL</b>	Indicates that the <i>cntr</i> value passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Getcncr**, **LAPI\_Waitcncr**

---

## LAPI\_Setcncr\_wstatus Subroutine

### Purpose

Used to set a counter to a specified value and to set the associated destination list array and destination status array to the counter.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Setcncr_wstatus(hndl, cncr, num_dest, dest_list, dest_status)
```

```
lapi_handle_t hndl;  
lapi_cncr_t *cncr;  
int num_dest;  
uint *dest_list;  
int *dest_status;
```

### FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_SETCNTR_WSTATUS(hndl, cncr, num_dest, dest_list, dest_status, ierror)  
INTEGER hndl  
TYPE (LAPI_CNTR_T) :: cncr  
INTEGER num_dest  
INTEGER dest_list(*)  
INTEGER dest_status  
INTEGER ierror
```

### Description

**Type of call:** recovery

This subroutine sets *cncr* to **0**. Use **LAPI\_Setcncr\_wstatus** to set the associated destination list array (*dest\_list*) and destination status array (*dest\_status*) to the counter. Use a corresponding **LAPI\_Nopoll\_wait** call to access these arrays. These arrays record the status of a task from where the thread calling **LAPI\_Nopoll\_wait()** is waiting for a response.

The return values for *dest\_status* are:

<b>LAPI_MSG_INITIAL</b>	The task is purged or is not received.
<b>LAPI_MSG_RECVD</b>	The task is received.
<b>LAPI_MSG_PURGED</b>	The task is purged, but not received.

**LAPI\_MSG\_PURGED\_RCVD** The task is received and then purged.

**LAPI\_MSG\_INVALID** Not valid; the task is already purged.

**Note:** To use this subroutine, the *lib\_vers* field in the **lapi\_info\_t** structure must be set to **L2\_LIB** or **LAST\_LIB**.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*num\_dest* Specifies the number of tasks in the destination list.

*dest\_list* Specifies an array of destinations waiting for this counter update. If the value of this parameter is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN), no status is returned to the user.

### INPUT/OUTPUT

*cntr* Specifies the address of the counter to be set (in C) or the counter structure (in FORTRAN). The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

### OUTPUT

*dest\_status* Specifies an array of status that corresponds to *dest\_list*. The value of this parameter can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

Use of this subroutine is *not* recommended on a system that is running Parallel Environment (PE).

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_CNTR\_NULL** Indicates that the *cntr* value passed in is NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

**LAPI\_ERR\_RET\_PTR\_NULL** Indicates that the value of *dest\_status* is NULL in C (or **LAPI\_ADDR\_NULL** in FORTRAN), but the value of *dest\_list* is not NULL in C (or **LAPI\_ADDR\_NULL** in FORTRAN).

## Location

`/usr/lib/liblapi_r.a`

## Related Information

Subroutines: **LAPI\_Getcntr**, **LAPI\_Nopoll\_wait**, **LAPI\_Purge\_totask**, **LAPI\_Setcntr**

---

## LAPI\_Term Subroutine

### Purpose

Terminates and cleans up a LAPI context.

## Library

Availability Library (**liblapi\_r.a**)

## C Syntax

```
#include <lapi.h>

int LAPI_Term(hdl)
lapi_handle_t hdl;
```

## FORTTRAN Syntax

```
include 'lapif.h'

LAPI_TERM(hdl, ierror)
INTEGER hdl
INTEGER ierror
```

## Description

**Type of call:** local termination

Use this subroutine to terminate the LAPI context that is specified by *hdl*. Any LAPI notification threads that are associated with this context are terminated. An error occurs when any LAPI calls are made using *hdl* after **LAPI\_Term** is called.

A DGSP that is registered under that LAPI handle remains valid even after **LAPI\_Term** is called on *hdl*.

## Parameters

### INPUT

*hdl* Specifies the LAPI handle.

### OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

LAPI statistics are *not* reported for shared memory communication and data transfer, or for messages that a task sends to itself.

## C Examples

To terminate a LAPI context (represented by *hdl*):

```
LAPI_Term(hdl);
```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hdl* passed in is not valid (not initialized or in terminated state).

## Location

**/usr/lib/liblapi\_r.a**

## Related Information

Subroutines: **LAPI\_Init**, **LAPI\_Purge\_totask**, **LAPI\_Resume\_totask**



---

## LAPI\_Util Subroutine

### Purpose

Serves as a wrapper function for such data gather/scatter operations as registration and reservation, for updating UDP port information, and for obtaining pointers to locking and signaling functions that are associated with a shared LAPI lock.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>

int LAPI_Util(hndl, util_cmd)
lapi_handle_t hndl;
lapi_util_t *util_cmd;
```

### FORTTRAN Syntax

```
include 'lapif.h'

LAPI_UTIL(hndl, util_cmd, ierror)
INTEGER hndl
TYPE (LAPI_UTIL_T) :: util_cmd
INTEGER ierror
```

### Description

**Type of call:** Data gather/scatter program (DGSP), UDP port information, and lock sharing utilities

This subroutine is used for several different operations, which are indicated by the command type value in the beginning of the command structure. The **lapi\_util\_t** structure is defined as:

```
typedef union {
    lapi_util_type_t    Util_type;
    lapi_reg_dgsp_t    RegDgsp;
    lapi_dref_dgsp_t    DrefDgsp;
    lapi_resv_dgsp_t    ResvDgsp;
    lapi_reg_ddm_t      DdmFunc;
    lapi_add_udp_port_t Udp;
    lapi_pack_dgsp_t    PackDgsp;
    lapi_unpack_dgsp_t  UnpackDgsp;
    lapi_thread_func_t  ThreadFunc;
} lapi_util_t;
```

The enumerated type **lapi\_util\_type\_t** has these values:

Table 1. *lapi\_util\_type\_t* types

Value of <i>Util_type</i>	Union member as interpreted by LAPI_Util
LAPI_REGISTER_DGSP	lapi_reg_dgsp_t
LAPI_UNRESERVE_DGSP	lapi_dref_dgsp_t
LAPI_RESERVE_DGSP	lapi_resv_dgsp_t
LAPI_REG_DDM_FUNC	lapi_reg_ddm_t
LAPI_ADD_UDP_DEST_PORT	lapi_add_udp_port_t
LAPI_DGSP_PACK	lapi_pack_dgsp_t
LAPI_DGSP_UNPACK	lapi_unpack_dgsp_t

Table 1. *lapi\_util\_type\_t* types (continued)

Value of <i>Util_type</i>	Union member as interpreted by <b>LAPI_Util</b>
<b>LAPI_GET_THREAD_FUNC</b>	<b>lapi_thread_func_t</b>

*hndl* is not checked for command type **LAPI\_REGISTER\_DGSP**, **LAPI\_RESERVE\_DGSP**, or **LAPI\_UNRESERVE\_DGSP**.

## **LAPI\_REGISTER\_DGSP**

You can use this operation to register a LAPI DGSP that you have created. To register a LAPI DGSP, **lapi\_dgsp\_descr\_t** *idgsp* must be passed in. LAPI returns a handle (**lapi\_dg\_handle\_t** *dgsp\_handle*) to use for all future LAPI calls. The *dgsp\_handle* that is returned by a register operation is identified as a **lapi\_dg\_handle\_t** type, which is the appropriate type for **LAPI\_Xfer** and **LAPI\_Util** calls that take a DGSP. This returned *dgsp\_handle* is also defined to be castable to a pointer to a **lapi\_dgsp\_descr\_t** for those situations where the LAPI user requires read-only access to information that is contained in the cached DGSP. The register operation delivers a DGSP to LAPI for use in future message send, receive, pack, and unpack operations. LAPI creates its own copy of the DGSP and protects it by reference count. All internal LAPI operations that depend on a DGSP cached in LAPI ensure the preservation of the DGSP by incrementing the reference count when they begin a dependency on the DGSP and decrementing the count when that dependency ends. A DGSP, once registered, can be used from any LAPI instance. **LAPI\_Term** does not discard any DGSPs.

You can register a DGSP, start one or more LAPI operations using the DGSP, and then unreserve it with no concern about when the LAPI operations that depend on the DGSP will be done using it. See **LAPI\_RESERVE\_DGSP** and **LAPI\_UNRESERVE\_DGSP** for more information.

In general, the DGSP you create and pass in to the **LAPI\_REGISTER\_DGSP** call using the *dgsp* parameter is discarded after LAPI makes and caches its own copy. Because DGSP creation is complex, user errors may occur, but extensive error checking at data transfer time would hurt performance. When developing code that creates DGSPs, you can invoke extra validation at the point of registration by setting the **LAPI\_VERIFY\_DGSP** environment variable. **LAPI\_Util** will return any detected errors. Any errors that exist and are not detected at registration time will cause problems during data transfer. Any errors detected during data transfer will be reported by an asynchronous error handler. A segmentation fault is one common symptom of a faulty DGSP. If multiple DGSPs are in use, the asynchronous error handler will not be able to identify which DGSP caused the error. For more information about asynchronous error handling, see **LAPI\_Init**.

**LAPI\_REGISTER\_DGSP** uses the **lapi\_reg\_dgsp\_t** command structure.

Table 2. The *lapi\_reg\_dgsp\_t* fields

<b>lapi_reg_dgsp_t</b> field	<b>lapi_reg_dgsp_t</b> field type	<b>lapi_reg_dgsp_t</b> usage
<i>Util_type</i>	<b>lapi_util_type_t</b>	<b>LAPI_REGISTER_DGSP</b>
<i>idgsp</i>	<b>lapi_dgsp_descr_t</b>	IN - pointer to DGSP program
<i>dgsp_handle</i>	<b>lapi_dg_handle_t</b>	OUT - handle for a registered DGSP program
<i>in_usr_func</i>	<b>lapi_usr_fcall_t</b>	For debugging only
<i>status</i>	<b>lapi_status_t</b>	OUT - future support

## **LAPI\_RESERVE\_DGSP**

You can use this operation to reserve a DGSP. This operation is provided because a LAPI client might cache a LAPI DGSP handle for later use. The client needs to ensure the DGSP will not be discarded before the cached handle is used. A DGSP handle, which is defined to be a pointer to a DGSP description that is already cached inside LAPI, is passed to this operation. The DGSP handle is also defined to be a structure pointer, so that client programs can get direct access to information in the DGSP. Unless the client can be certain that the DGSP will not be "unreserved" by another thread while it is being accessed,

the client should bracket the access window with its own reserve/unreserve operation. The client is not to modify the cached DGSP, but LAPI has no way to enforce this. The reserve operation increments the user reference count, thus protecting the DGSP until an unreserve operation occurs. This is needed because the thread that placed the reservation will expect to be able to use or examine the cached DGSP until it makes an unreserve call (which decrements the user reference count), even if the unreserve operation that matches the original register operation occurs within this window on some other thread.

**LAPI\_RESERVE\_DGSP** uses the **lapi\_resv\_dgsp\_t** command structure.

Table 3. The *lapi\_resv\_dgsp\_t* fields

<b>lapi_resv_dgsp_t field</b>	<b>lapi_resv_dgsp_t field type</b>	<b>lapi_resv_dgsp_t usage</b>
<i>Util_type</i>	<b>lapi_util_type_t</b>	<b>LAPI_RESERVE_DGSP</b>
<i>dgsp_handle</i>	<b>lapi_dg_handle_t</b>	OUT - handle for a registered DGSP program
<i>in_usr_func</i>	<b>lapi_usr_fcall_t</b>	For debugging only
<i>status</i>	<b>lapi_status_t</b>	OUT - future support

## LAPI\_UNRESERVE\_DGSP

You can use this operation to unregister or unreserve a DGSP. This operation decrements the user reference count. If external and internal reference counts are zero, this operation lets LAPI free the DGSP. All operations that decrement a reference count cause LAPI to check to see if the counts have both become 0 and if they have, dispose of the DGSP. Several internal LAPI activities increment and decrement a second reference count. The cached DGSP is disposable only when all activities (internal and external) that depend on it and use reference counting to preserve it have discharged their reference. The DGSP handle is passed to LAPI as a value parameter and LAPI does not nullify the caller's handle. It is your responsibility to not use this handle again because in doing an unreserve operation, you have indicated that you no longer count on the handle remaining valid.

**LAPI\_UNRESERVE\_DGSP** uses the **lapi\_dref\_dgsp\_t** command structure.

Table 4. The *lapi\_dref\_dgsp\_t* fields

<b>lapi_dref_dgsp_t field</b>	<b>lapi_dref_dgsp_t field type</b>	<b>lapi_dref_dgsp_t usage</b>
<i>Util_type</i>	<b>lapi_util_type_t</b>	<b>LAPI_UNRESERVE_DGSP</b>
<i>dgsp_handle</i>	<b>lapi_dg_handle_t</b>	OUT - handle for a registered DGSP program
<i>in_usr_func</i>	<b>lapi_usr_fcall_t</b>	For debugging only
<i>status</i>	<b>lapi_status_t</b>	OUT - future support

## LAPI\_REG\_DDM\_FUNC

You can use this operation to register data distribution manager (DDM) functions. It works in conjunction with the DGSM CONTROL instruction. Primarily, it is used for **MPI\_Accumulate**, but LAPI clients can provide any DDM function. It is also used to establish a callback function for processing data that is being scattered into a user buffer on the destination side.

The native LAPI user can install a callback without affecting the one MPI has registered for **MPI\_Accumulate**. The function prototype for the callback function is:

```
typedef long ddm_func_t (          /* return number of bytes processed */
    void *in,                    /* pointer to inbound data          */
    void *inout,                 /* pointer to destination space     */
    long bytes,                  /* number of bytes inbound         */
    int operand,                 /* CONTROL operand value           */
    int operation                 /* CONTROL operation value         */
);
```

A DDM function acts between the arrival of message data and the target buffer. The most common usage is to combine inbound data with data already in the target buffer. For example, if the target buffer is an array of integers and the incoming message consists of integers, the DDM function can be written to add each incoming integer to the value that is already in the buffer. The *operand* and *operation* fields of the DDM function allow one DDM function to support a range of operations with the CONTROL instruction by providing the appropriate values for these fields.

See *RSCT for AIX 5L: LAPI Programming Guide* for more information about DGSP programming.

**LAPI\_REG\_DDM\_FUNC** uses the `lapi_reg_ddm_t` command structure. Each call replaces the previous function pointer, if there was one.

Table 5. The `lapi_reg_ddm_t` fields

<code>lapi_reg_ddm_t</code> field	<code>lapi_reg_ddm_t</code> field type	<code>lapi_reg_ddm_t</code> usage
<code>Util_type</code>	<code>lapi_util_type_t</code>	<b>LAPI_REG_DDM_FUNC</b>
<code>ddm_func</code>	<code>ddm_func_t *</code>	IN - DDM function pointer
<code>in_usr_func</code>	<code>lapi_usr_fcall_t</code>	For debugging only
<code>status</code>	<code>lapi_status_t</code>	OUT - future support

## LAPI\_DGSP\_PACK

You can use this operation to gather data to a pack buffer from a user buffer under control of a DGSP. A single buffer may be packed by a series of calls. The caller provides a *position* value that is initialized to the starting offset within the buffer. Each pack operation adjusts *position*, so the next pack operation can begin where the previous pack operation ended. In general, a series of pack operations begins with *position* initialized to 0, but any offset is valid. There is no state carried from one pack operation to the next. Each pack operation starts at the beginning of the DGSP it is passed.

**LAPI\_DGSP\_PACK** uses the `lapi_pack_dgsp_t` command structure.

Table 6. The `lapi_pack_dgsp_t` fields

<code>lapi_pack_dgsp_t</code> field	<code>lapi_pack_dgsp_t</code> field type	<code>lapi_pack_dgsp_t</code> usage
<code>Util_type</code>	<code>lapi_util_type_t</code>	<b>LAPI_DGSP_PACK</b>
<code>dgsp_handle</code>	<code>lapi_dg_handle_t</code>	OUT - handle for a registered DGSP program
<code>in_buf</code>	<code>void *</code>	IN - source buffer to pack
<code>bytes</code>	<code>ulong</code>	IN - number of bytes to pack
<code>out_buf</code>	<code>void *</code>	OUT - output buffer for pack
<code>out_size</code>	<code>ulong</code>	IN - output buffer size in bytes
<code>position</code>	<code>ulong</code>	IN/OUT - current buffer offset
<code>in_usr_func</code>	<code>lapi_usr_fcall_t</code>	For debugging only
<code>status</code>	<code>lapi_status_t</code>	OUT - future support

## LAPI\_DGSP\_UNPACK

You can use this operation to scatter data from a packed buffer to a user buffer under control of a DGSP. A single buffer may be unpacked by a series of calls. The caller provides a *position* value that is initialized to the starting offset within the packed buffer. Each unpack operation adjusts *position*, so the next unpack operation can begin where the previous unpack operation ended. In general, a series of unpack operations begins with *position* initialized to 0, but any offset is valid. There is no state carried from one unpack operation to the next. Each unpack operation starts at the beginning of the DGSP it is passed.

**LAPI\_DGSP\_UNPACK** uses the **lapi\_unpack\_dgsp\_t** command structure.

Table 7. The *lapi\_unpack\_dgsp\_t* fields

<b>lapi_unpack_dgsp_t field</b>	<b>lapi_unpack_dgsp_t field type</b>	<b>lapi_unpack_dgsp_t usage</b>
<i>Util_type</i>	<b>lapi_util_type_t</b>	<b>LAPI_DGSP_UNPACK</b>
<i>dgsp_handle</i>	<b>lapi_dg_handle_t</b>	OUT - handle for a registered DGSP program
<i>buf</i>	<b>void *</b>	IN - source buffer for unpack
<i>in_size</i>	<b>ulong</b>	IN - source buffer size in bytes
<i>out_buf</i>	<b>void *</b>	OUT - output buffer for unpack
<i>bytes</i>	<b>ulong</b>	IN - number of bytes to unpack
<i>out_size</i>	<b>ulong</b>	IN - output buffer size in bytes
<i>position</i>	<b>ulong</b>	IN/OUT - current buffer offset
<i>in_usr_func</i>	<b>lapi_usr_fcall_t</b>	For debugging only
<i>status</i>	<b>lapi_status_t</b>	OUT - future support

## LAPI\_ADD\_UDP\_DEST\_PORT

You can use this operation to update UDP port information about the destination task. This operation can be used when you have written your own UDP handler (*udp\_hdlr*) and you need to support recovery of failed tasks. You cannot use this operation under the POE runtime environment.

**LAPI\_ADD\_UDP\_DEST\_PORT** uses the **lapi\_add\_udp\_port\_t** command structure.

Table 8. The *lapi\_add\_udp\_port\_t* fields

<b>lapi_add_udp_port_t field</b>	<b>lapi_add_udp_port_t field type</b>	<b>lapi_add_udp_port_t usage</b>
<i>Util_type</i>	<b>lapi_util_type_t</b>	<b>LAPI_ADD_UDP_DEST_PORT</b>
<i>tgt</i>	<b>uint</b>	IN - destination task ID
<i>udp_port</i>	<b>lapi_udp_t *</b>	IN - UDP port information for the target
<i>instance_no</i>	<b>uint</b>	IN - Instance number of UDP
<i>in_usr_func</i>	<b>lapi_usr_fcall_t</b>	For debugging only
<i>status</i>	<b>lapi_status_t</b>	OUT - future support

## LAPI\_GET\_THREAD\_FUNC

You can use this operation to retrieve various shared locking and signalling functions. Retrieval of these functions is valid only after LAPI is initialized and before LAPI is terminated. You should not call any of these functions after LAPI is terminated.

**LAPI\_GET\_THREAD\_FUNC** uses the **lapi\_thread\_func\_t** command structure.

Table 9. The *lapi\_thread\_func\_t* fields

<b>lapi_thread_func_t field</b>	<b>lapi_thread_func_t field type</b>	<b>lapi_thread_func_t usage</b>
<i>Util_type</i>	<b>lapi_util_type_t</b>	<b>LAPI_GET_THREAD_FUNC</b>
<i>mutex_lock</i>	<b>lapi_mutex_lock_t</b>	OUT - mutex lock function pointer
<i>mutex_unlock</i>	<b>lapi_mutex_unlock_t</b>	OUT - mutex unlock function pointer
<i>mutex_trylock</i>	<b>lapi_mutex_trylock_t</b>	OUT - mutex try lock function pointer
<i>mutex_getowner</i>	<b>lapi_mutex_getowner_t</b>	OUT - mutex get owner function pointer
<i>cond_wait</i>	<b>lapi_cond_wait_t</b>	OUT - condition wait function pointer
<i>cond_timedwait</i>	<b>lapi_cond_timedwait_t</b>	OUT - condition timed wait function pointer

Table 9. The `lapi_thread_func_t` fields (continued)

<code>lapi_thread_func_t</code> field	<code>lapi_thread_func_t</code> field type	<code>lapi_thread_func_t</code> usage
<code>cond_signal</code>	<code>lapi_cond_signal_t</code>	OUT - condition signal function pointer
<code>cond_init</code>	<code>lapi_cond_init_t</code>	OUT - initialize condition function pointer
<code>cond_destroy</code>	<code>lapi_cond_destroy_t</code>	OUT - destroy condition function pointer

LAPI uses the pthread library for thread ID management. You can therefore use `pthread_self()` to get the running thread ID and `lapi_mutex_getowner_t` to get the thread ID that owns the shared lock. Then, you can use `pthread_equal()` to see if the two are the same.

**Mutex thread functions:** `LAPI_GET_THREAD_FUNC` includes the following mutex thread functions: mutex lock, mutex unlock, mutex try lock, and mutex get owner.

### Mutex lock function pointer

```
int (*lapi_mutex_lock_t)(lapi_handle_t hndl);
```

This function acquires the lock that is associated with the specified LAPI handle. The call blocks if the lock is already held by another thread. Deadlock can occur if the calling thread is already holding the lock. You are responsible for preventing and detecting deadlocks.

### Parameters

#### INPUT

`hndl` Specifies the LAPI handle.

### Return values

**0** Indicates that the lock was acquired successfully.

**EINVAL** Is returned if the lock is not valid because of an incorrect `hndl` value.

### Mutex unlock function pointer

```
int (*lapi_mutex_unlock_t)(lapi_handle_t hndl);
```

This function releases the lock that is associated with the specified LAPI handle. A thread should only unlock its own locks.

### Parameters

#### INPUT

`hndl` Specifies the LAPI handle.

### Return values

**0** Indicates that the lock was released successfully.

**EINVAL** Is returned if the lock is not valid because of an incorrect `hndl` value.

### Mutex try lock function pointer

```
int (*lapi_mutex_trylock_t)(lapi_handle_t hndl);
```

This function tries to acquire the lock that is associated with the specified LAPI handle, but returns immediately if the lock is already held.

### Parameters

## INPUT

*hdl* Specifies the LAPI handle.

## Return values

**0** Indicates that the lock was acquired successfully.

**EBUSY** Indicates that the lock is being held.

**EINVAL** Is returned if the lock is not valid because of an incorrect *hdl* value.

## Mutex get owner function pointer

```
int (*lapi_mutex_getowner_t)(lapi_handle_t hndl, pthread_t *tid);
```

This function gets the pthread ID of the thread that is currently holding the lock associated with the specified LAPI handle. **LAPI\_NULL\_THREAD\_ID** indicates that the lock is not held at the time the function is called.

## Parameters

### INPUT

*hdl* Specifies the LAPI handle.

### OUTPUT

*tid* Is a pointer to hold the pthread ID to be retrieved.

## Return values

**0** Indicates that the lock owner was retrieved successfully.

**EINVAL** Is returned if the lock is not valid because of an incorrect *hdl* value.

**Condition functions:** **LAPI\_GET\_THREAD\_FUNC** includes the following condition functions: condition wait, condition timed wait, condition signal, initialize condition, and destroy condition.

## Condition wait function pointer

```
int (*lapi_cond_wait_t)(lapi_handle_t hndl, lapi_cond_t *cond);
```

This function waits on a condition variable (*cond*). The user must hold the lock associated with the LAPI handle (*hdl*) before making the call. Upon the return of the call, LAPI guarantees that the lock is still being held. The same LAPI handle must be supplied to concurrent **lapi\_cond\_wait\_t** operations on the same condition variable.

## Parameters

### INPUT

*hdl* Specifies the LAPI handle.

*cond* Is a pointer to the condition variable to be waited on.

## Return values

**0** Indicates that the condition variable has been signaled.

**EINVAL** Indicates that the value specified by *hdl* or *cond* is not valid.

## Condition timed wait function pointer

```
int (*lapi_cond_timedwait_t)(lapi_handle_t hndl,
                            lapi_cond_t *cond,
                            struct timespec *timeout);
```

This function waits on a condition variable (*cond*). The user must hold the lock associated with the LAPI handle (*hndl*) before making the call. Upon the return of the call, LAPI guarantees that the lock is still being held. The same LAPI handle must be supplied to concurrent **lapi\_cond\_timedwait\_t** operations on the same condition variable.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*cond* Is a pointer to the condition variable to be waited on.

*timeout* Is a pointer to the absolute time structure specifying the timeout.

### Return values

**0** Indicates that the condition variable has been signaled.

**ETIMEDOUT** Indicates that time specified by *timeout* has passed.

**EINVAL** Indicates that the value specified by *hndl*, *cond*, or *timeout* is not valid.

### Condition signal function pointer

```
int (*lapi_cond_wait_t)(lapi_handle_t hndl, lapi_cond_t *cond);
typedef int (*lapi_cond_signal_t)(lapi_handle_t hndl, lapi_cond_t *cond);
```

This function signals a condition variable (*cond*) to wake up a thread that is blocked on the condition. If there are multiple threads waiting on the condition variable, which thread to wake up is decided randomly.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*cond* Is a pointer to the condition variable to be signaled.

### Return values

**0** Indicates that the condition variable has been signaled.

**EINVAL** Indicates that the value specified by *hndl* or *cond* is not valid.

### Initialize condition function pointer

```
int (*lapi_cond_init_t)(lapi_handle_t hndl, lapi_cond_t *cond);
```

This function initializes a condition variable.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.

*cond* Is a pointer to the condition variable to be initialized.

### Return values

**0** Indicates that the condition variable was initialized successfully.



<b>EAGAIN</b>	Indicates that the system lacked the necessary resources (other than memory) to initialize another condition variable.
<b>ENOMEM</b>	Indicates that there is not enough memory to initialize the condition variable.
<b>EINVAL</b>	Is returned if the <i>hdl</i> value is not valid.

### Destroy condition function pointer

```
int (*lapi_cond_destroy_t)(lapi_handle_t hndl, lapi_cond_t *cond);
```

This function destroys a condition variable.

### Parameters

#### INPUT

<i>hdl</i>	Specifies the LAPI handle.
<i>cond</i>	Is a pointer to the condition variable to be destroyed.

### Return values

<b>0</b>	Indicates that the condition variable was destroyed successfully.
<b>EBUSY</b>	Indicates that the implementation has detected an attempt to destroy the object referenced by <i>cond</i> while it is referenced (while being used in a <b>lapi_cond_wait_t</b> or <b>lapi_cond_timedwait_t</b> by another thread, for example).
<b>EINVAL</b>	Indicates that the value specified by <i>hdl</i> or <i>cond</i> is not valid.

### Parameters

#### INPUT

<i>hdl</i>	Specifies the LAPI handle.
------------	----------------------------

#### INPUT/OUTPUT

<i>util_cmd</i>	Specifies the command type of the utility function.
-----------------	---

#### OUTPUT

<i>ierror</i>	Specifies a FORTRAN return code. This is always the last parameter.
---------------	---

### Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_DGSP</b>	Indicates that the DGSP that was passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN) or is not a registered DGSP.
<b>LAPI_ERR_DGSP_ATOM</b>	Indicates that the DGSP has an <i>atom_size</i> that is less than <b>0</b> or greater than <b>MAX_ATOM_SIZE</b> .
<b>LAPI_ERR_DGSP_BRANCH</b>	Indicates that the DGSP attempted a branch that fell outside of the code array. This is returned only in validation mode.
<b>LAPI_ERR_DGSP_COPY_SZ</b>	Is returned with DGSP validation turned on when <b>MCOPY</b> block < <b>0</b> or <b>COPY</b> instruction with bytes < <b>0</b> . This is returned only in validation mode.
<b>LAPI_ERR_DGSP_FREE</b>	Indicates that LAPI tried to free a DGSP that is not valid or is no longer registered. There should be one <b>LAPI_UNRESERVE_DGSP</b> operation to

close the **LAPI\_REGISTER\_DGSP** operation and one **LAPI\_UNRESERVE\_DGSP** operation for each **LAPI\_RESERVE\_DGSP** operation.

- LAPI\_ERR\_DGSP\_OPC** Indicates that the DGSP *opcode* is not valid. This is returned only in validation mode.
- LAPI\_ERR\_DGSP\_STACK** Indicates that the DGSP has a greater GOSUB depth than the allocated stack supports. Stack allocation is specified by the *dgsp->depth* member. This is returned only in validation mode.
- LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).
- LAPI\_ERR\_MEMORY\_EXHAUSTED** Indicates that LAPI is unable to obtain memory from the system.
- LAPI\_ERR\_UDP\_PORT\_INFO** Indicates that the *udp\_port* information pointer is NULL (in C) or that the value of *udp\_port* is **LAPI\_ADDR\_NULL** (in FORTRAN).
- LAPI\_ERR\_UTIL\_CMD** Indicates that the command type is not valid.

## C Examples

1. To create and register a DGSP:

```
{
    /*
    ** DGSP code array. DGSP instructions are stored
    ** as ints (with constants defined in lapi.h for
    ** the number of integers needed to store each
    ** instruction). We will have one COPY and one ITERATE
    ** instruction in our DGSP. We use LAPI's constants
    ** to allocate the appropriate storage.
    */
    int code[LAPI_DGSM_COPY_SIZE+LAPI_DGSM_ITERATE_SIZE];

    /* DGSP description */
    lapi_dgsp_descr_t dgsp_d;

    /*
    ** Data structure for the xfer call.
    */
    lapi_xfer_t xfer_struct;

    /* DGSP data structures */
    lapi_dgsm_copy_t *copy_p; /* copy instruction */
    lapi_dgsm_iterate_t *iter_p; /* iterate instruction */
    int *code_ptr; /* code pointer */

    /* constant for holding code array info */
    int code_less_iterate_size;

    /* used for DGSP registration */
    lapi_reg_dgsp_t reg_util;

    /*
    ** Set up dgsp description
    */

    /* set pointer to code array */
    dgsp_d.code = &code[0];
}
```

```

/* set size of code array */
dgsp_d.code_size = LAPI_DGSM_COPY_SIZE + LAPI_DGSM_ITERATE_SIZE;

/* not using DGSP gosub instruction */
dgsp_d.depth = 1;

/*
** set density to show internal gaps in the
** DGSP data layout
*/
dgsp_d.density = LAPI_DGSM_SPARSE;

/* transfer 4 bytes at a time */
dgsp_d.size = 4;

/* advance the template by 8 for each iteration */
dgsp_d.extent = 8;

/*
** ext specifies the memory 'footprint' of
** data to be transferred. The lext specifies
** the offset from the base address to begin
** viewing the data. The rext specifies the
** length from the base address to use.
*/
dgsp_d.lext = 0;
dgsp_d.rext = 4;
/* atom size of 0 lets LAPI choose the packet size */
dgsp_d.atom_size = 0;

/*
** set up the copy instruction
*/
copy_p = (lapi_dgsm_copy_t *) (dgsp_d.code);
copy_p->opcode = LAPI_DGSM_COPY;

/* copy 4 bytes at a time */
copy_p->bytes = (long) 4;

/* start at offset 0 */
copy_p->offset = (long) 0;

/* set code pointer to address of iterate instruction */
code_less_iterate_size = dgsp_d.code_size - LAPI_DGSM_ITERATE_SIZE;
code_ptr = ((int *) (code)) + code_less_iterate_size;

/*
** Set up iterate instruction
*/
iter_p = (lapi_dgsm_iterate_t *) code_ptr;
iter_p->opcode = LAPI_DGSM_ITERATE;
iter_p->iter_loc = (-code_less_iterate_size);

/* Set up and do DGSP registration */
reg_util.Util_type = LAPI_REGISTER_DGSP;
reg_util.idgsp = &dgsp_d;
LAPI_Util(hndl, (lapi_util_t *) &reg_util);

/*
** LAPI returns a usable DGSP handle in
** reg_util.dgsp_handle
** Use this handle for subsequent reserve/unreserve
** and Xfer calls. On the receive side, this
** handle can be returned by the header handler using the
** return_info_t mechanism. The DGSP will then be used for

```

```
    ** scattering data.  
    */  
}
```

## 2. To reserve a DGSP handle:

```
{  
  
    reg_util.dgsp_handle = dgsp_handle;  
  
    /*  
    ** dgsp_handle has already been created and  
    ** registered as in the above example  
    */  
  
    reg_util.Util_type = LAPI_RESERVE_DGSP;  
    LAPI_Util(hndl, (lapi_util_t *)&reg_util);  
  
    /*  
    ** LAPI's internal reference count to dgsp_handle  
    ** will be incremented. DGSP will  
    ** remain available until an unreserve is  
    ** done for each reserve, plus one more for  
    ** the original registration.  
    */  
}
```

## 3. To unreserve a DGSP handle:

```
{  
  
    reg_util.dgsp_handle = dgsp_handle;  
  
    /*  
    ** dgsp_handle has already created and  
    ** registered as in the above example, and  
    ** this thread no longer needs it.  
    */  
  
    reg_util.Util_type = LAPI_UNRESERVE_DGSP;  
    LAPI_Util(hndl, (lapi_util_t *)&reg_util);  
  
    /*  
    ** An unreserve is required for each reserve,  
    ** plus one more for the original registration.  
    */  
}
```

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: [LAPI\\_Init](#), [LAPI\\_Xfer](#)

---

## LAPI\_Waitcnt Subroutine

### Purpose

Waits until a specified counter reaches the value specified.

## Library

Availability Library (**liblapi\_r.a**)

## C Syntax

```
#include <lapi.h>

int LAPI_Waitcntr(hndl, cntr, val, cur_cntr_val)
lapi_handle_t hndl;
lapi_cntr_t *cntr;
int val;
int *cur_cntr_val;
```

## FORTRAN Syntax

```
include 'lapif.h'

LAPI_WAITCNTR(hndl, cntr, val, cur_cntr_val, ierror)
INTEGER hndl
TYPE (LAPI_CNTR_T) :: cntr
INTEGER val
INTEGER cur_cntr_val
INTEGER ierror
```

## Description

**Type of call:** local progress monitor (blocking)

This subroutine waits until *cntr* reaches or exceeds the specified *val*. Once *cntr* reaches *val*, *cntr* is decremented by the value of *val*. In this case, "decremented" is used (as opposed to "set to zero") because *cntr* could have contained a value that was greater than the specified *val* when the call was made. This call may or may not check for message arrivals over the LAPI context *hndl*. The *cur\_cntr\_val* variable is set to the current counter value.

## Parameters

### INPUT

*hndl* Specifies the LAPI handle.  
*val* Specifies the value the counter needs to reach.

### INPUT/OUTPUT

*cntr* Specifies the counter structure (in FORTRAN) to be waited on or its address (in C). The value of this parameter cannot be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

### OUTPUT

*cur\_cntr\_val* Specifies the integer value of the current counter. This value can be NULL (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).  
*ierror* Specifies a FORTRAN return code. This is always the last parameter.

## Restrictions

LAPI statistics are *not* reported for shared memory communication and data transfer, or for messages that a task sends to itself.

## C Examples

To wait on a counter to reach a specified value:

```

{
    int          val;
    int          cur_cntr_val;
    lapi_cntr_t  some_cntr;
    .
    .
    .
    LAPI_Waitcptr(hndl, &some_cntr, val, &cur_cntr_val);
    /* Upon return, some_cntr has reached val */
}

```

## Return Values

**LAPI\_SUCCESS** Indicates that the function call completed successfully.

**LAPI\_ERR\_CNTR\_NULL** Indicates that the *cntr* pointer is NULL (in C) or that the value of *cntr* is **LAPI\_ADDR\_NULL** (in FORTRAN).

**LAPI\_ERR\_HNDL\_INVALID** Indicates that the *hndl* passed in is not valid (not initialized or in terminated state).

## Location

/usr/lib/liblapi\_r.a

## Related Information

Subroutines: **LAPI\_Amsend**, **LAPI\_Amsendv**, **LAPI\_Get**, **LAPI\_Getcptr**, **LAPI\_Getv**, **LAPI\_Put**, **LAPI\_Putv**, **LAPI\_Rmw**, **LAPI\_Rmw64**, **LAPI\_Setcptr**, **LAPI\_Xfer**

---

## LAPI\_Xfer Subroutine

### Purpose

Serves as a wrapper function for LAPI data transfer functions.

### Library

Availability Library (**liblapi\_r.a**)

### C Syntax

```
#include <lapi.h>
```

```
int LAPI_Xfer(hndl, xfer_cmd)
lapi_handle_t hndl;
lapi_xfer_t *xfer_cmd;
```

```
typedef struct {
    uint          src;          /* Target task ID */
    uint          reason;      /* LAPI return codes */
    ulong         reserve[6];  /* Reserved */
} lapi_sh_info_t;
```

```
typedef void (scompl_hdlr_t)(lapi_handle_t *hndl, void *completion_param,
                             lapi_sh_info_t *info);
```

## FORTRAN Syntax

```
include 'lapif.h'
```

```
LAPI_XFER(hdl, xfer_cmd, ierror)  
INTEGER hdl  
TYPE (fortran_xfer_type) :: xfer_cmd  
INTEGER ierror
```

## Description

**Type of call:** point-to-point communication (non-blocking)

The **LAPI\_Xfer** subroutine provides a superset of the functionality of these subroutines: **LAPI\_Amsend**, **LAPI\_Amsendv**, **LAPI\_Put**, **LAPI\_Putv**, **LAPI\_Get**, **LAPI\_Getv**, and **LAPI\_Rmw**. In addition, **LAPI\_Xfer** provides data gather/scatter program (DGSP) messages transfer.

In C, the **LAPI\_Xfer** command is passed a pointer to a union. It examines the first member of the union, **Xfer\_type**, to determine the transfer type, and to determine which union member was passed. **LAPI\_Xfer** expects every field of the identified union member to be set. It does not examine or modify any memory outside of the identified union member. **LAPI\_Xfer** treats all union members (except **status**) as read-only data.

This subroutine provides the following functions:

- The remote address fields are expanded to be of type **lapi\_long\_t**, which is long enough for a 64-bit address. This allows a 32-bit task to send data to 64-bit addresses, which may be important in client/server programs.
- **LAPI\_Xfer** allows the origin counter to be replaced with a send completion callback.
- **LAPI\_Xfer** is used to transfer data using LAPI's data gather/scatter program (DGSP) interface.

The **lapi\_xfer\_t** structure is defined as:

```
typedef union {  
    lapi_xfer_type_t    Xfer_type;  
    lapi_get_t          Get;  
    lapi_am_t           Am;  
    lapi_rmw_t          Rmw;  
    lapi_put_t          Put;  
    lapi_getv_t         Getv;  
    lapi_putv_t         Putv;  
    lapi_amv_t          Amv;  
    lapi_amdgspt        Dgsp;  
} lapi_xfer_t;
```

Though the **lapi\_xfer\_t** structure applies only to the C version of **LAPI\_Xfer**, the following tables include the FORTRAN equivalents of the C datatypes.

Table 10 list the values of the **lapi\_xfer\_type\_t** structure for C and the explicit *Xfer\_type* values for FORTRAN.

Table 10. LAPI\_Xfer structure types

Value of <i>Xfer_type</i> (C or FORTRAN)	Union member as interpreted by LAPI_Xfer (C)	Value of <i>fortran_xfer_type</i> (FORTRAN)
LAPI_AM_XFER	<b>lapi_am_t</b>	LAPI_AM_T
LAPI_AMV_XFER	<b>lapi_amv_t</b>	LAPI_AMV_T
LAPI_DGSP_XFER	<b>lapi_amdgspt</b>	LAPI_AMDGSP_T
LAPI_GET_XFER	<b>lapi_get_t</b>	LAPI_GET_T
LAPI_GETV_XFER	<b>lapi_getv_t</b>	LAPI_GETV_T

Table 10. LAPI\_Xfer structure types (continued)

Value of <i>Xfer_type</i> (C or FORTRAN)	Union member as interpreted by LAPI_Xfer (C)	Value of <i>fortran_xfer_type</i> (FORTRAN)
LAPI_PUT_XFER	<code>lapi_put_t</code>	LAPI_PUT_T
LAPI_PUTV_XFER	<code>lapi_putv_t</code>	LAPI_PUTV_T
LAPI_RMW_XFER	<code>lapi_rmw_t</code>	LAPI_RMW_T

## lapi\_am\_t details

Table 11 shows the correspondence among the parameters of the **LAPI\_Amsend** subroutine, the fields of the C **lapi\_am\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_am\_t** fields are listed in Table 11 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 11. LAPI\_Amsend and lapi\_am\_t equivalents

lapi_am_t field name (C)	lapi_am_t field type (C)	Equivalent FORTRAN datatype	Equivalent LAPI_Amsend parameter
<i>Xfer_type</i>	<code>lapi_xfer_type_t</code>	INTEGER(KIND = 4)	implicit in C  LAPI_Xfer value in FORTRAN: LAPI_AM_XFER
<i>flags</i>	<code>int</code>	INTEGER(KIND = 4)	none  LAPI_Xfer parameter in FORTRAN: <i>flags</i>
<i>tgt</i>	<code>uint</code>	INTEGER(KIND = 4)	<i>tgt</i>
none	none	INTEGER(KIND = 4)	LAPI_Xfer parameter in FORTRAN: <i>pad</i>
<i>hdr_hdl</i>	<code>lapi_long_t</code>	INTEGER(KIND = 8)	<i>hdr_hdl</i>
<i>uhdr_len</i>	<code>uint</code>	INTEGER(KIND = 4)	<i>uhdr_len</i>
none	none	INTEGER(KIND = 4)	LAPI_Xfer parameter in FORTRAN (64-bit): <i>pad2</i>
<i>uhdr</i>	<code>void *</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	<i>uhdr</i>
<i>udata</i>	<code>void *</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	<i>udata</i>
<i>udata_len</i>	<code>ulong</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	<i>udata_len</i>
<i>shdlr</i>	<code>scompl_hndlr_t *</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	none  LAPI_Xfer parameter in FORTRAN: <i>shdlr</i>
<i>sinfo</i>	<code>void *</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	none  LAPI_Xfer parameter in FORTRAN: <i>sinfo</i>
<i>tgt_cntr</i>	<code>lapi_long_t</code>	INTEGER(KIND = 8)	<i>tgt_cntr</i>
<i>org_cntr</i>	<code>lapi_cntr_t *</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	<i>org_cntr</i>
<i>cmpl_cntr</i>	<code>lapi_cntr_t *</code>	INTEGER(KIND = 4) (32-bit) INTEGER(KIND = 8) (64-bit)	<i>cmpl_cntr</i>

When the origin data buffer is free to be used, the pointer to the send completion handler (*shdlr*) is called



with the send completion data (*sinfo*) if *shdlr* is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). Otherwise, the behavior is identical to that of **LAPI\_Amsend**.

### **lapi\_amv\_t details**

Table 12 shows the correspondence among the parameters of the **LAPI\_Amsendv** subroutine, the fields of the C **lapi\_amv\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_amv\_t** fields are listed in Table 12 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 12. *LAPI\_Amsendv* and *lapi\_amv\_t* equivalents

<b>lapi_amv_t field name (C)</b>	<b>lapi_amv_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Amsendv parameter</b>
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	<b>INTEGER(KIND = 4)</b>	implicit in C  <b>LAPI_Xfer</b> value in FORTRAN: <b>LAPI_AMV_XFER</b>
<i>flags</i>	<b>int</b>	<b>INTEGER(KIND = 4)</b>	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>flags</i>
<i>tgt</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	<i>tgt</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN: <i>pad</i>
<i>hdr_hdl</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>hdr_hdl</i>
<i>uhdr_len</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	<i>uhdr_len</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (64-bit): <i>pad2</i>
<i>uhdr</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>uhdr</i>
<i>shdlr</i>	<b>scompl_hndlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>shdlr</i>
<i>sinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>sinfo</i>
<i>org_vec</i>	<b>lapi_vec_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_vec</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (32-bit): <i>pad2</i>
<i>tgt_cntr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_cntr</i>
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_cntr</i>
<i>cmpl_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>cmpl_cntr</i>

### **lapi\_amdgspl\_t details**

Table 13 on page 674 shows the correspondence among the fields of the C **lapi\_amdgspl\_t** structure and their datatypes, how they are used in **LAPI\_Xfer**, and the equivalent FORTRAN datatypes. The **lapi\_amdgspl\_t** fields are listed in Table 13 on page 674 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 13. The *lapi\_amdgsp\_t* fields

<b>lapi_amdgsp_t field name (C)</b>	<b>lapi_amdgsp_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>LAPI_Xfer usage</b>
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	<b>INTEGER(KIND = 4)</b>	<b>LAPI_DGSP_XFER</b>
<i>flags</i>	<b>int</b>	<b>INTEGER(KIND = 4)</b>	This field allows users to specify directives or hints to LAPI. If you do not want to use any directives or hints, you must set this field to <b>0</b> . See “The <i>lapi_amdgsp_t</i> flags field” for more information.
<i>tgt</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	target task
none	none	<b>INTEGER(KIND = 4)</b>	<i>pad</i> (padding alignment for FORTRAN only)
<i>hdr_hdl</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	header handler to invoke at target
<i>uhdr_len</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	user header length (multiple of processor’s doubleword size)
none	none	<b>INTEGER(KIND = 4)</b>	<i>pad2</i> (padding alignment for 64-bit FORTRAN only)
<i>uhdr</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	pointer to user header
<i>udata</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	pointer to user data
<i>udata_len</i>	<b>ulong</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	user data length
<i>shdlr</i>	<b>scompl_hdlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	send completion handler (optional)
<i>sinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	data pointer to pass to send completion handler (optional)
<i>tgt_cntr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	target counter (optional)
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	origin counter (optional)
<i>cmpl_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	completion counter (optional)
<i>dgsp</i>	<b>lapi_dg_handle_t</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	Handle of a registered DGSP
<i>status</i>	<b>lapi_status_t</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	Status to return (future use)
none	none	<b>INTEGER(KIND = 4)</b>	<i>pad3</i> (padding alignment for 64-bit FORTRAN only)

When the origin data buffer is free to be modified, the send completion handler (*shdlr*) is called with the send completion data (*sinfo*), if *shdlr* is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN).

See “Using *lapi\_am\_dgsp\_t* for scatter-side DGSP processing” on page 675 for more information.

**The *lapi\_amdgsp\_t* flags field:** One or more flags can be set using the **|** (bitwise or) operator. User directives are always followed and could result in incorrect results if used improperly. Appropriate hints might improve performance, but they may be ignored by LAPI. Inappropriate hints might degrade performance, but they will not cause incorrect results.

The following directive flags are defined:

**USE\_TGT\_VEC\_TYPE** Instructs LAPI to use the vector type of the target vector (*tgt\_vec*). In other words, *tgt\_vec* is to be interpreted as type **lapi\_vec\_t**; otherwise, it is interpreted as type **lapi\_lvec\_t**. The **lapi\_lvec\_t** type uses **lapi\_long\_t**. The **lapi\_vec\_t** type uses **void \*** or **long**. Incorrect results will occur if one type is used in place of the other.

**BUFFER\_BOTH\_CONTIGUOUS**

Instructs LAPI to treat all data to be transferred as contiguous, which can improve performance. If this flag is set when non-contiguous data is sent, data will likely be corrupted.

The following hint flags are defined:

**LAPI\_NOT\_USE\_BULK\_XFER**

Instructs LAPI not to use bulk transfer, independent of the current setting for the task.

**LAPI\_USE\_BULK\_XFER**

Instructs LAPI to use bulk transfer, independent of the current setting for the task.

If neither of these hint flags is set, LAPI will use the behavior defined for the task. If both of these hint flags are set, **LAPI\_NOT\_USE\_BULK\_XFER** will take precedence.

These hints may or may not be honored by the communication library.

**Using lapi\_am\_dgsp\_t for scatter-side DGSP processing:** Beginning with AIX 5.2, LAPI allows additional information to be returned from the header handler through the use of the **lapi\_return\_info\_t** datatype. See *RSCT for AIX 5L: LAPI Programming Guide* for more information about **lapi\_return\_info\_t**. In the case of transfer type **lapi\_am\_dgsp\_t**, this mechanism can be used to instruct LAPI to run a user DGSP to scatter data on the receive side.

To use this mechanism, pass a **lapi\_return\_info\_t \*** pointer back to LAPI through the *msg\_len* member of the user header handler. The *dgsp\_handle* member of the passed structure must point to a DGSP description that has been registered on the receive side. See **LAPI\_Util** and *RSCT for AIX 5L: LAPI Programming Guide* for details on building and registering DGSPs.

**lapi\_get\_t details**

Table 14 shows the correspondence among the parameters of the **LAPI\_Get** subroutine, the fields of the C **lapi\_get\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_get\_t** fields are listed in Table 14 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 14. *LAPI\_Get* and *lapi\_get\_t* equivalents

<b>lapi_get_t</b> field name (C)	<b>lapi_get_t</b> field type (C)	Equivalent FORTRAN datatype	Equivalent <b>LAPI_Get</b> parameter
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	INTEGER(KIND = 4)	implicit in C  <b>LAPI_Xfer</b> value in FORTRAN: <b>LAPI_GET_XFER</b>
<i>flags</i>	<b>int</b>	INTEGER(KIND = 4)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>flags</i>
<i>tgt</i>	<b>uint</b>	INTEGER(KIND = 4)	<i>tgt</i>
none	none	INTEGER(KIND = 4)	<b>LAPI_Xfer</b> parameter in FORTRAN: <i>pad</i>

Table 14. *LAPI\_Get* and *lapi\_get\_t* equivalents (continued)

<b>lapi_get_t field name (C)</b>	<b>lapi_get_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Get parameter</b>
<i>tgt_addr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_addr</i>
<i>org_addr</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_addr</i>
<i>len</i>	<b>ulong</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>len</i>
<i>tgt_cntr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_cntr</i>
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_cntr</i>
<i>chndlr</i>	<b>compl_hndlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none <b>LAPI_Xfer</b> parameter in FORTRAN: <i>chndlr</i>
<i>cinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none <b>LAPI_Xfer</b> parameter in FORTRAN: <i>cinfo</i>

When the origin data buffer has completely arrived, the pointer to the completion handler (*chndlr*) is called with the completion data (*cinfo*), if *chndlr* is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). Otherwise, the behavior is identical to that of **LAPI\_Get**.

### **lapi\_getv\_t details**

Table 15 shows the correspondence among the parameters of the **LAPI\_Getv** subroutine, the fields of the C **lapi\_getv\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_getv\_t** fields are listed in Table 14 on page 675 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 15. *LAPI\_Getv* and *lapi\_getv\_t* equivalents

<b>lapi_getv_t field name (C)</b>	<b>lapi_getv_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Getv parameter</b>
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	<b>INTEGER(KIND = 4)</b>	implicit in C <b>LAPI_Xfer</b> value in FORTRAN: <b>LAPI_GETV_XFER</b>
<i>flags</i>	<b>int</b>	<b>INTEGER(KIND = 4)</b>	none <b>LAPI_Xfer</b> parameter in FORTRAN: <i>flags</i>
<i>tgt</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	<i>tgt</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (64-bit): <i>pad</i>
<i>org_vec</i>	<b>lapi_vec_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_vec</i>
<i>tgt_vec</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>tgt_vec</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (32-bit): <i>pad</i>
<i>tgt_cntr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_cntr</i>
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_cntr</i>

Table 15. LAPI\_Getv and lapi\_getv\_t equivalents (continued)

<b>lapi_getv_t field name (C)</b>	<b>lapi_getv_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Getv parameter</b>
<i>chndlr</i>	<b>compl_hndlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>chndlr</i>
<i>cinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>cinfo</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (32-bit): <i>pad2</i>

The **flags** field accepts **USE\_TGT\_VEC\_TYPE** (see “The lapi\_amdgp\_t flags field” on page 674) to indicate that **tgt\_vec** is to be interpreted as type **lapi\_vec\_t**; otherwise, it is interpreted as type **lapi\_lvec\_t**. Note the corresponding field is **lapi\_vec\_t** in the **LAPI\_Getv** call.

When the origin data buffer has completely arrived, the pointer to the completion handler (**chndlr**) is called with the completion data (**cinfo**) if **chndlr** is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). Otherwise, the behavior is identical to that of **LAPI\_Getv**.

### lapi\_put\_t details

Table 16 shows the correspondence among the parameters of the **LAPI\_Put** subroutine, the fields of the C **lapi\_put\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_put\_t** fields are listed in Table 16 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 16. LAPI\_Put and lapi\_put\_t equivalents

<b>lapi_put_t field name (C)</b>	<b>lapi_put_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Put parameter</b>
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	<b>INTEGER(KIND = 4)</b>	implicit in C  <b>LAPI_Xfer</b> value in FORTRAN: <b>LAPI_PUT_XFER</b>
<i>flags</i>	<b>int</b>	<b>INTEGER(KIND = 4)</b>	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>flags</i>
<i>tgt</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	<i>tgt</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN: <i>pad</i>
<i>tgt_addr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_addr</i>
<i>org_addr</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_addr</i>
<i>len</i>	<b>ulong</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>len</i>
<i>shdlr</i>	<b>scompl_hndlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>shdlr</i>
<i>sinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>sinfo</i>

Table 16. LAPI\_Put and lapi\_put\_t equivalents (continued)

<b>lapi_put_t field name (C)</b>	<b>lapi_put_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Put parameter</b>
<i>tgt_cntr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_cntr</i>
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_cntr</i>
<i>cmpl_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>cmpl_cntr</i>

When the origin data buffer is free to be used, the pointer to the send completion handler (*shdlr*) is called with the send completion data (*sinfo*), if *shdlr* is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). Otherwise, the behavior is identical to that of **LAPI\_Put**.

### lapi\_putv\_t details

Table 17 shows the correspondence among the parameters of the **LAPI\_Putv** subroutine, the fields of the C **lapi\_putv\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_putv\_t** fields are listed in Table 16 on page 677 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 17. LAPI\_Putv and lapi\_putv\_t equivalents

<b>lapi_putv_t field name (C)</b>	<b>lapi_putv_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Putv parameter</b>
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	<b>INTEGER(KIND = 4)</b>	implicit in C  <b>LAPI_Xfer</b> value in FORTRAN: <b>LAPI_PUT_XFER</b>
<i>flags</i>	<b>int</b>	<b>INTEGER(KIND = 4)</b>	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>flags</i>
<i>tgt</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	<i>tgt</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (64-bit): <i>pad</i>
<i>shdlr</i>	<b>scompl_hndlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>shdlr</i>
<i>sinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>sinfo</i>
<i>org_vec</i>	<b>lapi_vec_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_vec</i>
<i>tgt_vec</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>tgt_vec</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (32-bit): <i>pad</i>
<i>tgt_cntr</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_cntr</i>
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_cntr</i>
<i>cmpl_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>cmpl_cntr</i>

The **flags** field accepts **USE\_TGT\_VEC\_TYPE** (see “The *lapi\_amdgp\_t* flags field” on page 674) to

indicate that **tgt\_vec** is to be interpreted as **lapi\_vec\_t**; otherwise, it is interpreted as a **lapi\_ivec\_t**. Note that the corresponding field is **lapi\_vec\_t** in the **LAPI\_Putv** call.

When the origin data buffer is free to be modified, the pointer to the send completion handler (**shdlr**) is called with the send completion data (**sinfo**), if **shdlr** is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). Otherwise, the behavior is identical to that of **LAPI\_Putv**.

## lapi\_rmw\_t details

Table 18 shows the correspondence among the parameters of the **LAPI\_Rmw** subroutine, the fields of the C **lapi\_rmw\_t** structure and their datatypes, and the equivalent FORTRAN datatypes. The **lapi\_rmw\_t** fields are listed in Table 16 on page 677 in the order that they occur in the **lapi\_xfer\_t** structure.

Table 18. *LAPI\_Rmw* and *lapi\_rmw\_t* equivalents

<b>lapi_rmw_t field name (C)</b>	<b>lapi_rmw_t field type (C)</b>	<b>Equivalent FORTRAN datatype</b>	<b>Equivalent LAPI_Rmw parameter</b>
<i>Xfer_type</i>	<b>lapi_xfer_type_t</b>	<b>INTEGER(KIND = 4)</b>	implicit in C  <b>LAPI_Xfer</b> value in FORTRAN: <b>LAPI_RMW_XFER</b>
<i>op</i>	<b>Rmw_ops_t</b>	<b>INTEGER(KIND = 4)</b>	<i>op</i>
<i>tgt</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	<i>tgt</i>
<i>size</i>	<b>uint</b>	<b>INTEGER(KIND = 4)</b>	implicit in C  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>size</i> (must be <b>32</b> or <b>64</b> )
<i>tgt_var</i>	<b>lapi_long_t</b>	<b>INTEGER(KIND = 8)</b>	<i>tgt_var</i>
<i>in_val</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>in_val</i>
<i>prev_tgt_val</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>prev_tgt_val</i>
<i>org_cntr</i>	<b>lapi_cntr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	<i>org_cntr</i>
<i>shdlr</i>	<b>scompl_hndlr_t *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>shdlr</i>
<i>sinfo</i>	<b>void *</b>	<b>INTEGER(KIND = 4)</b> (32-bit) <b>INTEGER(KIND = 8)</b> (64-bit)	none  <b>LAPI_Xfer</b> parameter in FORTRAN: <i>shdlr</i>
none	none	<b>INTEGER(KIND = 4)</b>	<b>LAPI_Xfer</b> parameter in FORTRAN (32-bit): <i>pad</i>

When the origin data buffer is free to be used, the pointer to the send completion handler (**shdlr**) is called with the send completion data (**sinfo**), if **shdlr** is not a NULL pointer (in C) or **LAPI\_ADDR\_NULL** (in FORTRAN). The *size* value must be either **32** or **64**, indicating whether you want the *in\_val* and *prev\_tgt\_val* fields to point to a 32-bit or 64-bit quantity, respectively. Otherwise, the behavior is identical to that of **LAPI\_Rmw**.

## Parameters

### INPUT

- hndl* Specifies the LAPI handle.
- xfer\_cmd* Specifies the name and parameters of the data transfer function.

## OUTPUT

*ierror* Specifies a FORTRAN return code. This is always the last parameter.

### Return Values

<b>LAPI_SUCCESS</b>	Indicates that the function call completed successfully.
<b>LAPI_ERR_DATA_LEN</b>	Indicates that the value of <i>udata_len</i> or <i>len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_DGSP</b>	Indicates that the DGSP that was passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN) or is not a registered DGSP.
<b>LAPI_ERR_DGSP_ATOM</b>	Indicates that the DGSP has an <i>atom_size</i> that is less than <b>0</b> or greater than <b>MAX_ATOM_SIZE</b> .
<b>LAPI_ERR_DGSP_BRANCH</b>	Indicates that the DGSP attempted a branch that fell outside the code array.
<b>LAPI_ERR_DGSP_CTL</b>	Indicates that a DGSP control instruction was encountered in a non-valid context (such as a gather-side control or scatter-side control with an atom size of <b>0</b> at gather, for example).
<b>LAPI_ERR_DGSP_OPC</b>	Indicates that the DGSP op-code is not valid.
<b>LAPI_ERR_DGSP_STACK</b>	Indicates that the DGSP has greater GOSUB depth than the allocated stack supports. Stack allocation is specified by the <i>dgsp-&gt;depth</i> member.
<b>LAPI_ERR_HDR_HNDLR_NULL</b>	Indicates that the <i>hdr_hdl</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_HNDL_INVALID</b>	Indicates that the <i>hdl</i> passed in is not valid (not initialized or in terminated state).
<b>LAPI_ERR_IN_VAL_NULL</b>	Indicates that the <i>in_val</i> pointer is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_MEMORY_EXHAUSTED</b>	LAPI is unable to obtain memory from the system.
<b>LAPI_ERR_OP_SZ</b>	Indicates that the <i>lapi_rmw_t size</i> field is not set to <b>32</b> or <b>64</b> .
<b>LAPI_ERR_ORG_ADDR_NULL</b>	Indicates either that the <i>udata</i> parameter passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN) and <i>udata_len</i> is greater than <b>0</b> , or that the <i>org_addr</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN) and <i>len</i> is greater than <b>0</b> .  <b>Note:</b> if <i>Xfer_type</i> = <b>LAPI_DGSP_XFER</b> , the case in which <i>udata</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN) and <i>udata_len</i> is greater than <b>0</b> is valid, so an error is not returned.
<b>LAPI_ERR_ORG_EXTENT</b>	Indicates that the <i>org_vec</i> 's extent ( <i>stride</i> * <i>num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_ORG_STRIDE</b>	Indicates that the <i>org_vec</i> stride is less than block.
<b>LAPI_ERR_ORG_VEC_ADDR</b>	Indicates that the <i>org_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but its length ( <i>org_vec-&gt;len[i]</i> ) is not <b>0</b> .
<b>LAPI_ERR_ORG_VEC_LEN</b>	Indicates that the sum of <i>org_vec-&gt;len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .



<b>LAPI_ERR_ORG_VEC_NULL</b>	Indicates that the <i>org_vec</i> value is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_ORG_VEC_TYPE</b>	Indicates that the <i>org_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_RMW_OP</b>	Indicates the op is not valid.
<b>LAPI_ERR_STRIDE_ORG_VEC_ADDR_NULL</b>	Indicates that the strided vector address <i>org_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_STRIDE_TGT_VEC_ADDR_NULL</b>	Indicates that the strided vector address <i>tgt_vec-&gt;info[0]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT</b>	Indicates that the <i>tgt</i> passed in is outside the range of tasks defined in the job.
<b>LAPI_ERR_TGT_ADDR_NULL</b>	Indicates that <i>ret_addr</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT_EXTENT</b>	Indicates that <i>tgt_vec</i> 's extent ( <i>stride</i> * <i>num_vecs</i> ) is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_TGT_PURGED</b>	Indicates that the subroutine returned early because <b>LAPI_Purge_totask()</b> was called.
<b>LAPI_ERR_TGT_STRIDE</b>	Indicates that the <i>tgt_vec</i> stride is less than block.
<b>LAPI_ERR_TGT_VAR_NULL</b>	Indicates that the <i>tgt_var</i> address is NULL (in C) or that the value of <i>tgt_var</i> is <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT_VEC_ADDR</b>	Indicates that the <i>tgt_vec-&gt;info[i]</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but its length ( <i>tgt_vec-&gt;len[i]</i> ) is not 0.
<b>LAPI_ERR_TGT_VEC_LEN</b>	Indicates that the sum of <i>tgt_vec-&gt;len</i> is greater than the value of LAPI constant <b>LAPI_MAX_MSG_SZ</b> .
<b>LAPI_ERR_TGT_VEC_NULL</b>	Indicates that <i>tgt_vec</i> is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN).
<b>LAPI_ERR_TGT_VEC_TYPE</b>	Indicates that the <i>tgt_vec-&gt;vec_type</i> is not valid.
<b>LAPI_ERR_UHDR_LEN</b>	Indicates that the <i>uhdr_len</i> value passed in is greater than <b>MAX_UHDR_SZ</b> or is not a multiple of the processor's doubleword size.
<b>LAPI_ERR_UHDR_NULL</b>	Indicates that the <i>uhdr</i> passed in is NULL (in C) or <b>LAPI_ADDR_NULL</b> (in FORTRAN), but <i>uhdr_len</i> is not 0.
<b>LAPI_ERR_VEC_LEN_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different lengths ( <i>len[ ]</i> ).
<b>LAPI_ERR_VEC_NUM_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different <i>num_vecs</i> .
<b>LAPI_ERR_VEC_TYPE_DIFF</b>	Indicates that <i>org_vec</i> and <i>tgt_vec</i> have different vector types ( <i>vec_type</i> ).
<b>LAPI_ERR_XFER_CMD</b>	Indicates that the <i>Xfer_cmd</i> is not valid.

## C Examples

1. To run the sample code shown in **LAPI\_Get** using the **LAPI\_Xfer** interface:

```

{
    lapi_xfer_t xfer_struct;

    /* initialize the table buffer for the data addresses */

```

```

/* get remote data buffer addresses */
LAPI_Address_init(hndl,(void *)data_buffer,data_buffer_list);
.
.
.
/* retrieve data_len bytes from address data_buffer_list[tgt] on */
/* task tgt. write the data starting at address data_buffer.    */
/* tgt_cntr and org_cntr can be NULL.                          */

xfer_struct.Get.Xfer_type = LAPI_GET_XFER;
xfer_struct.Get.flags = 0;
xfer_struct.Get.tgt = tgt;
xfer_struct.Get.tgt_addr = data_buffer_list[tgt];
xfer_struct.Get.org_addr = data_buffer;
xfer_struct.Get.len = data_len;
xfer_struct.Get.tgt_cntr = tgt_cntr;
xfer_struct.Get.org_cntr = org_cntr;

LAPI_Xfer(hndl, &xfer_struct);
}

```

2. To implement the **LAPI\_STRIDED\_VECTOR** example from **LAPI\_Amsendv** using the **LAPI\_Xfer** interface:

```

{
    lapi_xfer_t    xfer_struct;           /* info for LAPI_Xfer call */
    lapi_vec_t     vec;                  /* data for data transfer */
    .
    .
    .
    vec->num_vecs = NUM_VECS;           /* NUM_VECS = number of vectors to transfer */
                                        /* must match that of the target vector */
    vec->vec_type = LAPI_GEN_STRIDED_XFER; /* same as target vector */

    vec->info[0] = buffer_address;      /* starting address for data copy */
    vec->info[1] = block_size;          /* bytes of data to copy */
    vec->info[2] = stride;              /* distance from copy block to copy block */
    /* data will be copied as follows: */
    /* block_size bytes will be copied from buffer_address */
    /* block_size bytes will be copied from buffer_address+stride */
    /* block_size bytes will be copied from buffer_address+(2*stride) */
    /* block_size bytes will be copied from buffer_address+(3*stride) */
    .
    .
    .
    /* block_size bytes will be copied from buffer_address+((NUM_VECS-1)*stride) */
    .
    .
    .
    xfer_struct.Amv.Xfer_type = LAPI_AMV_XFER;
    xfer_struct.Amv.flags = 0;
    xfer_struct.Amv.tgt = tgt;
    xfer_struct.Amv.hdr_hdl = hdr_hdl_list[tgt];
    xfer_struct.Amv.uhdr_len = uhdr_len; /* user header length */
    xfer_struct.Amv.uhdr = uhdr;

    /* LAPI_AMV_XFER allows the use of a send completion handler */
    /* If non-null, the shdlr function is invoked at the point */
    /* the origin counter would increment. Note that both the */
    /* org_cntr and shdlr can be used. */
    /* The user's shdlr must be of type scompl_hdlr_t *. */
    /* scompl_hdlr_t is defined in /usr/include/lapi.h */
    xfer_struct.shdlr = shdlr;
}

```

```

    /* Use sinfo to pass user-defined data into the send      */
    /* completion handler, if desired.                       */
    xfer_struct.sinfo    = sinfo;        /* send completion data */

    xfer_struct.org_vec  = vec;
    xfer_struct.tgt_cntr = tgt_cntr;
    xfer_struct.org_cntr = org_cntr;
    xfer_struct.cmpl_cntr = cmpl_cntr;

    LAPI_Xfer(hndl, &xfer_struct);
    .
    .
}

```

See the **LAPI\_Amsendv** subroutine for more information about the header handler definition.

## Location

**/usr/lib/liblapi\_r.a**

## Related Information

Books: *RSCT for AIX 5L: LAPI Programming Guide* for information about bulk message transfer

Subroutines: **LAPI\_Amsend**, **LAPI\_Amsendv**, **LAPI\_Get**, **LAPI\_Getv**, **LAPI\_Put**, **LAPI\_Putv**, **LAPI\_Rmw**

## layout\_object\_create Subroutine

### Purpose

Initializes a layout context.

### Library

Layout Library (**libi18n.a**)

### Syntax

```
#include <sys/lc_layout.h>
```

```
int layout_object_create (locale_name, layout_object)
const char * locale_name;
LayoutObject * layout_object;
```

### Description

The **layout\_object\_create** subroutine creates the **LayoutObject** structure associated with the locale specified by the *locale\_name* parameter. The **LayoutObject** structure is a symbolic link containing all the data and methods necessary to perform the layout operations on context dependent and bidirectional characters of the locale specified.

When the **layout\_object\_create** subroutine completes without errors, the *layout\_object* parameter points to a valid **LayoutObject** structure that can be used by other BIDI subroutines. The returned **LayoutObject** structure is initialized to an initial state that defines the behavior of the BIDI subroutines. This initial state is locale dependent and is described by the layout values returned by the **layout\_object\_getvalue** subroutine. You can change the layout values of the **LayoutObject** structure using the **layout\_object\_setvalue** subroutine. Any state maintained by the **LayoutObject** structure is independent of the current global locale set with the **setlocale** subroutine.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the libcur Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

<i>locale_name</i>	Specifies a locale. It is recommended that you use the <b>LC_CTYPE</b> category by calling the <b>setlocale (LC_CTYPE, NULL)</b> subroutine.
<i>layout_object</i>	Points to a valid <b>LayoutObject</b> structure that can be used by other layout subroutines. This parameter is used only when the <b>layout_object_create</b> subroutine completes without errors.  The <i>layout_object</i> parameter is not set and a non-zero value is returned if a valid <b>LayoutObject</b> structure cannot be created.

## Return Values

Upon successful completion, the **layout\_object\_create** subroutine returns a value of 0. The *layout\_object* parameter points to a valid handle.

## Error Codes

If the **layout\_object\_create** subroutine fails, it returns the following error codes:

<b>LAYOUT_EINVAL</b>	The locale specified by the <i>locale_name</i> parameter is not available.
<b>LAYOUT_EMFILE</b>	The OPEN_MAX value of files descriptors are currently open in the calling process.
<b>LAYOUT_ENOMEM</b>	Insufficient storage space is available.

## Related Information

The “*layout\_object\_editshape* or *wcslayout\_object\_editshape* Subroutine,” “*layout\_object\_free* Subroutine” on page 695, “*layout\_object\_getvalue* Subroutine” on page 687, “*layout\_object\_setvalue* Subroutine” on page 689, “*layout\_object\_shapeboxchars* Subroutine” on page 691, “*layout\_object\_transform* or *wcslayout\_object\_transform* Subroutine” on page 692.

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## **layout\_object\_editshape** or **wcslayout\_object\_editshape** Subroutine

### Purpose

Edits the shape of the context text.

### Library

Layout library (**libi18n.a**)

## Syntax

```
#include <sys/lc_layout.h>
```

```
int layout_editshape ( layout_object, EditType, index, InpBuf, Inpsize, OutBuf, OutSize)
LayoutObject layout_object;
BooleanValue EditType;
size_t *index;
const char *InpBuf;
size_t *Inpsize;
void *OutBuf;
size_t *OutSize;

int wcslayout_object_editshape(layout_object, EditType, index, InpBuf, Inpsize, OutBuf, OutSize)
LayoutObject layout_object;
BooleanValue EditType;
size_t *index;
const wchar_t *InpBuf;
size_t *InpSize;
void *OutBuf;
size_t *OutSize;
```

## Description

The **layout\_object\_editshape** and **wcslayout\_object\_editshape** subroutines provide the shapes of the context text. The shapes are defined by the code element specified by the *index* parameter and any surrounding code elements specified by the ShapeContextSize layout value of the **LayoutObject** structure. The *layout\_object* parameter specifies this **LayoutObject** structure.

Use the **layout\_object\_editshape** subroutine when editing code elements of one byte. Use the **wcslayout\_object\_editshape** subroutine when editing single code elements of multibytes. These subroutines do not affect any state maintained by the **layout\_object\_transform** or **wcslayout\_object\_transform** subroutine.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the libcur Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

*layout\_object* Specifies the **LayoutObject** structure created by the **layout\_object\_create** subroutine.

<i>EditType</i>	<p>Specifies the type of edit shaping. When the <i>EditType</i> parameter stipulates the <i>EditInput</i> field, the subroutine reads the current code element defined by the <i>index</i> parameter and any preceding code elements defined by <i>ShapeContextSize</i> layout value of the <b>LayoutObject</b> structure. When the <i>EditType</i> parameter stipulates the <i>EditReplace</i> field, the subroutine reads the current code element defined by the <i>index</i> parameter and any surrounding code elements defined by <i>ShapeContextSize</i> layout value of the <b>LayoutObject</b> structure.</p> <p><b>Note:</b> The editing direction defined by the <i>Orientation</i> and <i>TEXT_VISUAL</i> of the <i>TypeOfText</i> layout values of the <b>LayoutObject</b> structure determines which code elements are preceding and succeeding.</p> <p>When the <i>ActiveShapeEditing</i> layout value of the <b>LayoutObject</b> structure is set to <i>True</i>, the <b>LayoutObject</b> structure maintains the state of the <i>EditInput</i> field that may affect subsequent calls to these subroutines with the <i>EditInput</i> field defined by the <i>EditType</i> parameter. The state of the <i>EditInput</i> field of <b>LayoutObject</b> structure is not affected when the <i>EditType</i> parameter is set to the <i>EditReplace</i> field. To reset the state of the <i>EditInput</i> field to its initial state, call these subroutines with the <i>InpBuf</i> parameter set to <i>NULL</i>. The state of the <i>EditInput</i> field is not affected if errors occur within the subroutines.</p>
<i>index</i>	<p>Specifies an offset (in bytes) to the start of a code element in the <i>InpBuf</i> parameter on input. The <i>InpBuf</i> parameter provides the base text to be edited. In addition, the context of the surrounding code elements is considered where the minimum set of code elements needed for the specific context dependent script(s) is identified by the <i>ShapeContextSize</i> layout value.</p> <p>If the set of surrounding code elements defined by the <i>index</i>, <i>InpBuf</i>, and <i>InpSize</i> parameters is less than the size of front and back of the <i>ShapeContextSize</i> layout value, these subroutines assume there is no additional context available. The caller must provide the minimum context if it is available. The <i>index</i> parameter is in units associated with the type of the <i>InpBuf</i> parameter.</p> <p>On return, the <i>index</i> parameter is modified to indicate the offset to the first code element of the <i>InpBuf</i> parameter that required shaping. The number of code elements that required shaping is indicated on return by the <i>InpSize</i> parameter.</p>
<i>InpBuf</i>	<p>Specifies the source to be processed. A <i>Null</i> value with the <i>EditInput</i> field in the <i>EditType</i> parameter indicates a request to reset the state of the <i>EditInput</i> field to its initial state.</p>
<i>InpSize</i>	<p>Any portion of the <i>InpBuf</i> parameter indicates the necessity for redrawing or shaping. Specifies the number of code elements to be processed in units on input. These units are associated with the types for these subroutines. A value of -1 indicates that the input is delimited by a <i>Null</i> code element.</p> <p>On return, the value is modified to the actual number of code elements needed by the <i>InpBuf</i> parameter. A value of 0 when the value of the <i>EditType</i> parameter is the <i>EditInput</i> field indicates that the state of the <i>EditInput</i> field is reset to its initial state. If the <i>OutBuf</i> parameter is not <i>NULL</i>, the respective shaped code elements are written into the <i>OutBuf</i> parameter.</p>
<i>OutBuf</i>	<p>Contains the shaped output text. You can specify this parameter as a <i>Null</i> pointer to indicate that no transformed text is required. If <i>Null</i>, the subroutines return the <i>index</i> and <i>InpSize</i> parameters, which specify the amount of text required, to be redrawn.</p> <p>The encoding of the <i>OutBuf</i> parameter depends on the <i>ShapeCharset</i> layout value defined in <i>layout_object</i> parameter. If the <i>ActiveShapeEditing</i> layout value is set to <i>False</i>, the encoding of the <i>OutBuf</i> parameter is to be the same as the code set of the locale associated with the specified <b>LayoutObject</b> structure.</p>

*OutSize* Specifies the size of the output buffer on input in number of bytes. Only the code elements required to be shaped are written into the *OutBuf* parameter.

The output buffer should be large enough to contain the shaped result; otherwise, only partial shaping is performed. If the *ActiveShapeEditing* layout value is set to True, the *OutBuf* parameter should be allocated to contain at least the number of code elements in the *InpBuf* parameter multiplied by the value of the *ShapeCharsetSize* layout value.

On return, the *OutSize* parameter is modified to the actual number of bytes placed in the output buffer.

When the *OutSize* parameter is specified as 0, the subroutines calculate the size of an output buffer large enough to contain the transformed text from the input buffer. The result will be returned in this field. The content of the buffers specified by the *InpBuf* and *OutBuf* parameters, and the value of the *InpSize* parameter, remain unchanged.

## Return Values

Upon successful completion, these subroutines return a value of 0. The *index* and *InpSize* parameters return the minimum set of code elements required to be redrawn.

## Error Codes

If these subroutines fail, they return the following error codes:

<b>LAYOUT_EILSEQ</b>	Shaping stopped due to an input code element that cannot be shaped. The <i>index</i> parameter indicates the code element that caused the error. This code element is either a valid code element that cannot be shaped according to the <i>ShapeCharset</i> layout value or an invalid code element not defined by the code set defined in the <b>LayoutObject</b> structure. Use the <b>mbtowc</b> or <b>wctomb</b> subroutine in the same locale as the <b>LayoutObject</b> structure to determine if the code element is valid.
<b>LAYOUT_E2BIG</b>	The output buffer is too small and the source text was not processed. The <i>index</i> and <i>InpSize</i> parameters are not guaranteed on return.
<b>LAYOUT_EINVAL</b>	Shaping stopped due to an incomplete code element or shift sequence at the end of input buffer. The <i>InpSize</i> parameter indicates the number of code elements successfully transformed. <b>Note:</b> You can use this error code to determine the code element causing the error.
<b>LAYOUT_ERANGE</b>	Either the <i>index</i> parameter is outside the range as defined by the <i>InpSize</i> parameter, more than 15 embedding levels are in the source text, or the <i>InpBuf</i> parameter contains unbalanced Directional Format (Push/Pop).

## Related Information

The “*layout\_object\_create* Subroutine” on page 683, “*layout\_object\_free* Subroutine” on page 695, “*layout\_object\_getvalue* Subroutine,” “*layout\_object\_setvalue* Subroutine” on page 689, “*layout\_object\_shapeboxchars* Subroutine” on page 691, “*layout\_object\_transform* or *wcslayout\_object\_transform* Subroutine” on page 692.

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## layout\_object\_getvalue Subroutine

### Purpose

Queries the current layout values of a **LayoutObject** structure.

## Library

Layout Library (**libi18n.a**)

## Syntax

```
#include <sys/lc_layout.h>
```

```
int layout_object_getvalue( layout_object, values, index)
LayoutObject layout_object;
LayoutValues values;
int *index;
```

## Description

The **layout\_object\_getvalue** subroutine queries the current setting of layout values within the **LayoutObject** structure. The *layout\_object* parameter specifies the **LayoutObject** structure created by the **layout\_object\_create** subroutine.

The name field of the **LayoutValues** structure contains the name of the layout value to be queried. The value field is a pointer to where the layout value is stored. The values are queried from the **LayoutObject** structure and represent its current state.

For example, if the layout value to be queried is of type T, the *value* parameter must be of type T\*. If T itself is a pointer, the **layout\_object\_getvalue** subroutine allocates space to store the actual data. The caller must free this data by calling the **free(T)** subroutine with the returned pointer.

When setting the value field, an extra level of indirection is present that is not present using the **layout\_object\_setvalue** parameter. When you set a layout value of type T, the value field contains T. However, when querying the same layout value, the value field contains &T.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the libcur Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

<i>layout_object</i>	Specifies the <b>LayoutObject</b> structure created by the <b>layout_object_create</b> subroutine.
<i>values</i>	Specifies an array of layout values of type <b>LayoutValueRec</b> that are to be queried in the <b>LayoutObject</b> structure. The end of the array is indicated by name=0.
<i>index</i>	Specifies a layout value to be queried. If the value cannot be queried, the <i>index</i> parameter causing the error is returned and the subroutine returns a non-zero value.

## Return Values

Upon successful completion, the **layout\_object\_getvalue** subroutine returns a value of 0. All layout values were successfully queried.

## Error Codes

If the **layout\_object\_getvalue** subroutine fails, it returns the following values:

<b>LAYOUT_EINVAL</b>	The layout value specified by the <i>index</i> parameter is unknown or the <i>layout_object</i> parameter is invalid.
<b>LAYOUT_EMOMEM</b>	Insufficient storage space is available.



## Examples

The following example queries whether the locale is bidirectional and gets the values of the in and out orientations.

```
#include <sys/lc_layout.h>
#include <locale.h>
main()
{
    LayoutObject plh;
    int RC=0;
    LayoutValues layout;
    LayoutTextDescriptor Descr;
    int index;

    RC=layout_object_create(setlocale(LC_CTYPE,""),&plh); /* create object */
    if (RC) {printf("Create error !!\n"); exit(0);}

    layout=malloc(3*sizeof(LayoutValueRec));
                                /* allocate layout array */
    layout[0].name=ActiveBidirection; /* set name */
    layout[1].name=Orientation;      /* set name */
    layout[1].value=(caddr_t)&Descr; /* send address of memory to be allocated by function */

    layout[2].name=0;              /* indicate end of array */
    RC=layout_object_getvalue(plh,layout,&index);
    if (RC) {printf("Getvalue error at %d !!\n",index); exit(0);}
    printf("ActiveBidirection = %d \n",*(layout[0].value));
                                /*print output*/
    printf("Orientation in = %x out = %x \n", Descr->>in, Descr->>out);

    free(layout);                /* free layout array */
    free (Descr);                /* free memory allocated by function */
    RC=layout_object_free(plh);  /* free layout object */
    if (RC) printf("Free error !!\n");
}
```

## Related Information

The “[layout\\_object\\_create Subroutine](#)” on page 683, “[layout\\_object\\_editshape or wcslayout\\_object\\_editshape Subroutine](#)” on page 684, “[layout\\_object\\_free Subroutine](#)” on page 695, “[layout\\_object\\_setvalue Subroutine](#),” “[layout\\_object\\_shapeboxchars Subroutine](#)” on page 691, and “[layout\\_object\\_transform or wcslayout\\_object\\_transform Subroutine](#)” on page 692.

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## layout\_object\_setvalue Subroutine

### Purpose

Sets the layout values of a **LayoutObject** structure.

### Library

Layout Library (**libi18n.a**)

### Syntax

```
#include <sys/lc_layout.h>
```

```
int layout_object_setvalue( layout_object, values, index)
LayoutObject layout_object;
LayoutValues values;
int *index;
```

## Description

The **layout\_object\_setvalue** subroutine changes the current layout values of the **LayoutObject** structure. The *layout\_object* parameter specifies the **LayoutObject** structure created by the **layout\_object\_create** subroutine. The values are written into the **LayoutObject** structure and may affect the behavior of subsequent layout functions.

**Note:** Some layout values do alter internal states maintained by a **LayoutObject** structure.

The name field of the LayoutValueRec structure contains the name of the layout value to be set. The value field contains the actual value to be set. The value field is large enough to support all types of layout values. For more information on layout value types, see "Layout Values for the Layout Library" in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the libcur Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

<i>layout_object</i>	Specifies the <b>LayoutObject</b> structure returned by the <b>layout_object_create</b> subroutine.
<i>values</i>	Specifies an array of layout values of the type LayoutValueRec that this subroutine sets. The end of the array is indicated by name=0.
<i>index</i>	Specifies a layout value to be queried. If the value cannot be queried, the index parameter causing the error is returned and the subroutine returns a non-zero value. If an error is generated, a subset of the values may have been previously set.

## Return Values

Upon successful completion, the **layout\_object\_setvalue** subroutine returns a value of 0. All layout values were successfully set.

## Error Codes

If the **layout\_object\_setvalue** subroutine fails, it returns the following values:

<b>LAYOUT_EINVAL</b>	The layout value specified by the <i>index</i> parameter is unknown, its value is invalid, or the <i>layout_object</i> parameter is invalid.
<b>LAYOUT_EMFILE</b>	The ( <b>OPEN_MAX</b> ) file descriptors are currently open in the calling process.
<b>LAYOUT_ENOMEM</b>	Insufficient storage space is available.

## Examples

The following example sets the TypeofText value to Implicit and the out value to Visual.

```
#include <sys/lc_layout.h>
#include <locale.h>

main()
{
LayoutObject plh;
int RC=0;
LayoutValues layout;
LayoutTextDescriptor Descr;
```

```

int index;

RC=layout_object_create(setlocale(LC_CTYPE,""),&plh); /* create object */
if (RC) {printf("Create error !!\n"); exit(0);}

layout=malloc(2*sizeof(LayoutValueRec)); /*allocate layout array*/
Descr=malloc(sizeof(LayoutTextDescriptorRec)); /* allocate text descriptor */
layout[0].name=TypeOfText;           /* set name */
layout[0].value=(caddr_t)Descr;      /* set value */
layout[1].name=0;                     /* indicate end of array */

Descr->in=TEXT_IMPLICIT;
Descr->out=TEXT_VISUAL; RC=layout_object_setvalue(plh,layout,&index);
if (RC) printf("SetValue error at %d!!\n",index); /* check return code */
free(layout);                          /* free allocated memory */
free (Descr);
RC=layout_object_free(plh);             /* free layout object */
if (RC) printf("Free error !!\n");
}

```

## Related Information

The “`layout_object_create` Subroutine” on page 683, “`layout_object_editshape` or `wcslayout_object_editshape` Subroutine” on page 684, “`layout_object_free` Subroutine” on page 695, “`layout_object_getvalue` Subroutine” on page 687, “`layout_object_shapeboxchars` Subroutine,” and “`layout_object_transform` or `wcslayout_object_transform` Subroutine” on page 692.

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## layout\_object\_shapeboxchars Subroutine

### Purpose

Shapes box characters.

### Library

Layout Library (`libi18n.a`)

### Syntax

```
#include <sys/lc_layout.h>
```

```

int layout_object_shapeboxchars( layout_object,  InpBuf,  InpSize,  OutBuf)
LayoutObject layout_object;
const char *InpBuf;
const size_t InpSize;
char *OutBuf;

```

### Description

The `layout_object_shapeboxchars` subroutine shapes box characters into the VT100 box character set.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the `libcure` Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

<i>layout_object</i>	Specifies the <b>LayoutObject</b> structure created by the <b>layout_object_create</b> subroutine.
<i>InpBuf</i>	Specifies the source text to be processed.
<i>InpSize</i>	Specifies the number of code elements to be processed.
<i>OutBuf</i>	Contains the shaped output text.

## Return Values

Upon successful completion, this subroutine returns a value of 0.

## Error Codes

If this subroutine fails, it returns the following values:

<b>LAYOUT_EILSEQ</b>	Shaping stopped due to an input code element that cannot be mapped into the VT100 box character set.
<b>LAYOUT_EINVAL</b>	Shaping stopped due to an incomplete code element or shift sequence at the end of the input buffer.

## Related Information

The “[layout\\_object\\_create Subroutine](#)” on page 683, “[layout\\_object\\_editshape or wcslayout\\_object\\_editshape Subroutine](#)” on page 684, “[layout\\_object\\_free Subroutine](#)” on page 695, “[layout\\_object\\_getvalue Subroutine](#)” on page 687, “[layout\\_object\\_setvalue Subroutine](#)” on page 689, and “[layout\\_object\\_transform or wcslayout\\_object\\_transform Subroutine](#).”

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## layout\_object\_transform or wcslayout\_object\_transform Subroutine

### Purpose

Transforms text according to the current layout values of a **LayoutObject** structure.

### Library

Layout Library (**libi18n.a**)

### Syntax

```
#include <sys/lc_layout.h>
int layout_object_transform ( layout_object, InpBuf, InpSize, OutBuf, OutSize, InpToOut, OutToInp, BidiLvl)
LayoutObject layout_object;
const char *InpBuf;
size_t *InpSize;
void * OutBuf;
size_t *OutSize;
size_t *InpToOut;
size_t *OutToInp;
unsigned char *BidiLvl;

int wcslayout_object_transform (layout_object, InpBuf, InpSize, OutBuf, OutSize, InpToOut, OutToInp, BidiLvl)
LayoutObject layout_object;
const char *InpBuf;
size_t *InpSize;
void *OutBuf;
```

```

Size_t *OutSize;
size_t *InpToOut;
size_t *OutToInp;
unsigned char *BidiLvl;

```

## Description

The **layout\_object\_transform** and **wcslayout\_object\_transform** subroutines transform the text specified by the *InpBuf* parameter according to the current layout values in the **LayoutObject** structure. Any layout value whose type is `LayoutTextDescriptor` describes the attributes within the *InpBuf* and *OutBuf* parameters. If the attributes are the same as the *InpBuf* and *OutBuf* parameters themselves, a null transformation is done with respect to that specific layout value.

The output of these subroutines may be one or more of the following results depending on the setting of the respective parameters:

<i>OutBuf</i> , <i>OutSize</i>	Any transformed data is stored in the <i>OutBuf</i> parameter.
<i>InpToOut</i>	A cross reference from each code element of the <i>InpBuf</i> parameter to the transformed data.
<i>OutToInp</i>	A cross reference to each code element of the <i>InpBuf</i> parameter from the transformed data.
<i>BidiLvl</i>	A weighted value that represents the directional level of each code element of the <i>InpBuf</i> parameter. The level is dependent on the internal directional algorithm of the <b>LayoutObject</b> structure.

You can specify each of these output parameters as `Null` to indicate that no output is needed for the specific parameter. However, you should set at least one of these parameters to a nonNULL value to perform any significant work.

To perform shaping of a text string without reordering of code elements, set the `TypeOfText` layout value to **TEXT\_VISUAL** and the in and out values of the `Orientation` layout value alike. These layout values are in the **LayoutObject** structure.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the `libcur` Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

<i>layout_object</i>	Specifies the <b>LayoutObject</b> structure created by the <b>layout_object_create</b> subroutine.
<i>InpBuf</i>	Specifies the source text to be processed. This parameter cannot be null.
<i>InpSize</i>	Specifies the units of code elements processed associated with the bytes for the <b>layout_object_transform</b> and <b>wcslayout_object_transform</b> subroutines. A value of -1 indicates that the input is delimited by a null code element. On return, the value is modified to the actual number of code elements processed in the <i>InpBuf</i> parameter. However, if the value in the <i>OutSize</i> parameter is zero, the value of the <i>InpSize</i> parameter is not changed.
<i>OutBuf</i>	Contains the transformed data. You can specify this parameter as a null pointer to indicate that no transformed data is required.

The encoding of the *OutBuf* parameter depends on the `ShapeCharset` layout value defined in the **LayoutObject** structure. If the `ActiveShapeEditing` layout value is set to `True`, the encoding of the *OutBuf* parameter is the same as the code set of the locale associated with the **LayoutObject** structure.

<i>OutSize</i>	<p>Specifies the size of the output buffer in number of bytes. The output buffer should be large enough to contain the transformed result; otherwise, only a partial transformation is performed. If the <code>ActiveShapeEditing</code> layout value is set to <code>True</code>, the <code>OutBuf</code> parameter should be allocated to contain at least the number of code elements multiplied by the <code>ShapeCharsetSize</code> layout value.</p> <p>On return, the <code>OutSize</code> parameter is modified to the actual number of bytes placed in this parameter.</p> <p>When you specify the <code>OutSize</code> parameter as 0, the subroutine calculates the size of an output buffer to be large enough to contain the transformed text. The result is returned in this field. The content of the buffers specified by the <code>InpBuf</code> and <code>OutBuf</code> parameters, and a value of the <code>InpSize</code> parameter remains unchanged.</p>
<i>InpToOut</i>	<p>Represents an array of values with the same number of code elements as the <code>InpBuf</code> parameter if <code>InpToOut</code> parameter is not a null pointer.</p> <p>On output, the <i>n</i>th value in <code>InpToOut</code> parameter corresponds to the <i>n</i>th code element in <code>InpBuf</code> parameter. This value is the index in <code>OutBuf</code> parameter which identifies the transformed <code>ShapeCharset</code> element of the <i>n</i>th code element in <code>InpBuf</code> parameter. You can specify the <code>InpToOut</code> parameter as null if no index array from the <code>InpBuf</code> to <code>OutBuf</code> parameters is desired.</p>
<i>OutToInp</i>	<p>Represents an array of values with the same number of code elements as contained in the <code>OutBuf</code> parameter if the <code>OutToInp</code> parameter is not a null pointer.</p> <p>On output, the <i>n</i>th value in the <code>OutToInp</code> parameter corresponds to the <i>n</i>th <code>ShapeCharset</code> element in the <code>OutBuf</code> parameter. This value is the index in the <code>InpBuf</code> parameter which identifies the original code element of the <i>n</i>th <code>ShapeCharset</code> element in the <code>OutBuf</code> parameter. You can specify the <code>OutToInp</code> parameter as <code>NULL</code> if no index array from the <code>OutBuf</code> to <code>InpBuf</code> parameters is desired.</p>
<i>BidiLvl</i>	<p>Represents an array of values with the same number of elements as the source text if the <code>BidiLvl</code> parameter is not a null pointer. The <i>n</i>th value in the <code>BidiLvl</code> parameter corresponds to the <i>n</i>th code element in the <code>InpBuf</code> parameter. This value is the level of this code element as determined by the bidirectional algorithm. You can specify the <code>BidiLvl</code> parameter as null if a levels array is not desired.</p>

## Return Values

Upon successful completion, these subroutines return a value of 0.

## Error Codes

If these subroutines fail, they return the following values:

<b>LAYOUT_EILSEQ</b>	<p>Transformation stopped due to an input code element that cannot be shaped or is invalid. The <code>InpSize</code> parameter indicates the number of the code element successfully transformed.</p> <p><b>Note:</b> You can use this error code to determine the code element causing the error.</p> <p>This code element is either a valid code element but cannot be shaped into the <code>ShapeCharset</code> layout value or is an invalid code element not defined by the code set of the locale of the <code>LayoutObject</code> structure. You can use the <code>mbtowc</code> and <code>wctomb</code> subroutines to determine if the code element is valid when used in the same locale as the <code>LayoutObject</code> structure.</p>
<b>LAYOUT_E2BIG</b>	<p>The output buffer is full and the source text is not entirely processed.</p>

## LAYOUT\_EINVAL

Transformation stopped due to an incomplete code element or shift sequence at the end of the input buffer. The *InpSize* parameter indicates the number of the code elements successfully transformed.

**Note:** You can use this error code to determine the code element causing the error.

## LAYOUT\_ERANGE

More than 15 embedding levels are in the source text or the *InpBuf* parameter contains unbalanced Directional Format (Push/Pop).

When the size of *OutBuf* parameter is not large enough to contain the entire transformed text, the input text state at the end of the **LAYOUT\_E2BIG** error code is returned. To resume the transformation on the remaining text, the application calls the **layout\_object\_transform** subroutine with the same **LayoutObject** structure, the same *InpBuf* parameter, and *InpSize* parameter set to 0.

## Examples

The following is an example of transformation of both directional re-ordering and shaping.

### Notes:

1. Uppercase represent left-to-right characters; lowercase represent right-to-left characters.
2. xyz represent the shapes of cde.

Position: 0123456789  
InpBuf: AB cde 12Z

Position: 0123456789  
OutBuf: AB 12 zyxZ

Position: 0123456789  
ToTarget: 0128765349

Position: 0123456789  
ToSource: 0127865439

Position: 0123456789  
BidiLevel: 0001111220

## Related Information

The “[layout\\_object\\_create Subroutine](#)” on page 683, “[layout\\_object\\_editshape](#) or [wcslayout\\_object\\_editshape Subroutine](#)” on page 684, “[layout\\_object\\_free Subroutine](#),” “[layout\\_object\\_getvalue Subroutine](#)” on page 687, “[layout\\_object\\_setvalue Subroutine](#)” on page 689, and “[layout\\_object\\_shapeboxchars Subroutine](#)” on page 691.

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## layout\_object\_free Subroutine

### Purpose

Frees a **LayoutObject** structure.

### Library

Layout library (**libi18n.a**)

### Syntax

```
#include <sys/lc_layout.h>
```

```
int layout_object_free(layout_object)
LayoutObject layout_object;
```

## Description

The **layout\_object\_free** subroutine releases all the resources of the **LayoutObject** structure created by the **layout\_object\_create** subroutine. The *layout\_object* parameter specifies this **LayoutObject** structure.

**Note:** If you are developing internationalized applications that may support multibyte locales, please see Use of the libcur Package in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*

## Parameters

*layout\_object* Specifies a **LayoutObject** structure returned by the **layout\_object\_create** subroutine.

## Return Values

Upon successful completion, the **layout\_object\_free** subroutine returns a value of 0. All resources associated with the *layout\_object* parameter are successfully deallocated.

## Error Codes

If the **layout\_object\_free** subroutine fails, it returns the following error code:

**LAYOUT\_EFAULT** Errors occurred while processing the request.

## Related Information

The “[layout\\_object\\_create Subroutine](#)” on page 683, “[layout\\_object\\_editshape or wcslayout\\_object\\_editshape Subroutine](#)” on page 684, “[layout\\_object\\_getvalue Subroutine](#)” on page 687, “[layout\\_object\\_setvalue Subroutine](#)” on page 689, “[layout\\_object\\_shapeboxchars Subroutine](#)” on page 691, and “[layout\\_object\\_transform or wcslayout\\_object\\_transform Subroutine](#)” on page 692.

Bidirectionality and Character Shaping and National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## Idahread Subroutine

### Purpose

Reads the archive header of a member of an archive file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ar.h>
#include <ldfcn.h>
```



```
int ldahread( ldPointer, ArchiveHeader)
LDFILE *ldPointer;
ARCHDR *ArchiveHeader;
```

## Description

If the **TYPE(ldPointer)** macro from the **ldfcn.h** file is the archive file magic number, the **ldahread** subroutine reads the archive header of the extended common object file currently associated with the *ldPointer* parameter into the area of memory beginning at the *ArchiveHeader* parameter.

## Parameters

<i>ldPointer</i>	Points to the <b>LDFILE</b> structure that was returned as the result of a successful call to <b>ldopen</b> or <b>ldaopen</b> .
<i>ArchiveHeader</i>	Points to a <b>ARCHDR</b> structure.

## Return Values

The **ldahread** subroutine returns a SUCCESS or FAILURE value.

## Error Codes

The **ldahread** routine fails if the **TYPE(ldPointer)** macro does not represent an archive file, or if it cannot read the archive header.

## Related Information

The **ldfhread** (“ldfhread Subroutine” on page 701) subroutine, **ldgetname** (“ldgetname Subroutine” on page 702) subroutine, **ldlread**, **ldlinit**, or **ldlitem** (“ldlread, ldlinit, or ldlitem Subroutine” on page 704) subroutine, **ldshread** or **ldnshread** (“ldshread or ldnshread Subroutine” on page 710) subroutine, **ldtbread** (“ldtbread Subroutine” on page 714) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldclose or ldaclose Subroutine

### Purpose

Closes a common object file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldclose( ldPointer)
LDFILE *ldPointer;
```

```
int ldaclose(ldPointer)
LDFILE *ldPointer;
```

## Description

The **ldopen** and **ldclose** subroutines provide uniform access to both simple object files and object files that are members of archive files. Thus, an archive of common object files can be processed as if it were a series of simple common object files.

If the **ldfcn.h** file **TYPE(*ldPointer*)** macro is the magic number of an archive file, and if there are any more files in the archive, the **ldclose** subroutine reinitializes the **ldfcn.h** file **OFFSET(*ldPointer*)** macro to the file address of the next archive member and returns a failure value. The **ldfile** structure is prepared for a subsequent **ldopen**.

If the **TYPE(*ldPointer*)** macro does not represent an archive file, the **ldclose** subroutine closes the file and frees the memory allocated to the **ldfile** structure associated with *ldPointer*.

The **ldaclose** subroutine closes the file and frees the memory allocated to the **ldfile** structure associated with the *ldPointer* parameter regardless of the value of the **TYPE(*ldPointer*)** macro.

## Parameters

*ldPointer*            Pointer to the **LDFILE** structure that was returned as the result of a successful call to **ldopen** or **ldaopen**.

## Return Values

The **ldclose** subroutine returns a SUCCESS or FAILURE value.

The **ldaclose** subroutine always returns a SUCCESS value and is often used in conjunction with the **ldaopen** subroutine.

## Error Codes

The **ldclose** subroutine returns a failure value if there are more files to archive.

## Related Information

The **ldaopen** or **ldopen** (“ldopen or ldaopen Subroutine” on page 707) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldexpd32, ldexpd64, and ldexpd128 Subroutines

### Purpose

Loads exponent of a decimal floating-point number.

### Syntax

```
#include <math.h>
```

```
_Decimal132 ldexpd32 (x, exp)  
_Decimal132 x;  
int exp;
```

```
_Decimal164 ldexpd64 (x, exp)  
_Decimal164 x;  
int exp;
```

```
_Decimal128 ldexpd128 (x, exp)
_Decimal128 x;
int exp;
```

## Description

The **ldexpd32**, **ldexpd64**, and **ldexpd128** subroutines compute the quantity  $x * 10^{exp}$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.  
*exp* Specifies the exponent of 10.

## Return Values

Upon successful completion, the **ldexpd32**, **ldexpd64**, and **ldexpd128** subroutines return *x* multiplied by 10, raised to the power *exp*.

If the **ldexpd32**, **ldexpd64**, or **ldexpd128** subroutines would cause overflow, a range error occurs and the **ldexpd32**, **ldexpd64**, and **ldexpd128** subroutines return  $\pm$ **HUGE\_VAL\_D32**,  $\pm$ **HUGE\_VAL\_D64**, and  $\pm$ **HUGE\_VAL\_D128** (according to the sign of *x*), respectively.

If the correct value would cause underflow, and is not representable, a range error may occur, and 0.0 is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or Inf, *x* is returned.

If *exp* is 0, *x* is returned.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

## Related Information

“**feclearexcept** Subroutine” on page 265, “**fetestexcept** Subroutine” on page 273, and “**class**, **\_class**, **finite**, **isnan**, or **unordered** Subroutines” on page 167

**math.h** in in *AIX 5L Version 5.3 Files Reference*.

---

## ldexp, ldexpf, or ldexpl Subroutine

### Purpose

Loads exponent of a floating-point number.

### Syntax

```
#include <math.h>
float ldexpf (x, exp)
float x;
```

```
int exp;
```

```
long double ldexpl (x, exp)  
long double x;  
int exp;
```

```
double ldexp (x, exp)  
double x;  
int exp;
```

## Description

The **ldexpf**, **ldexpl**, and **ldexp** subroutines compute the quantity  $x * 2^{exp}$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.  
*exp* Specifies the exponent of 2.

## Return Values

Upon successful completion, the **ldexpf**, **ldexpl**, and **ldexp** subroutines return *x* multiplied by 2, raised to the power *exp*.

If the **ldexpf**, **ldexpl**, or **ldexp** subroutines would cause overflow, a range error occurs and the **ldexpf**, **ldexpl**, and **ldexp** subroutines return **±HUGE\_VALF**, **±HUGE\_VALL**, and **±HUGE\_VAL** (according to the sign of *x*), respectively.

If the correct value would cause underflow, and is not representable, a range error may occur, and 0.0 is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or Inf, *x* is returned.

If *exp* is 0, *x* is returned.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

## Error Codes

If the result of the **ldexp** or **ldexpl** subroutine overflows, then **±HUGE\_VAL** is returned, and the global variable **errno** is set to **ERANGE**.

If the result of the **ldexp** or **ldexpl** subroutine underflows, 0 is returned, and the **errno** global variable is set to a **ERANGE** value.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167

## ldfhread Subroutine

### Purpose

Reads the file header of an XCOFF file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldfhread ( ldPointer, FileHeader)
LDFILE *ldPointer;
void *FileHeader;
```

### Description

The **ldfhread** subroutine reads the file header of the object file currently associated with the *ldPointer* parameter into the area of memory beginning at the *FileHeader* parameter. For AIX 4.3.2 and above, it is the responsibility of the calling routine to provide a pointer to a buffer large enough to contain the file header of the associated object file. Since the **ldopen** subroutine provides magic number information (via the **HEADER(ldPointer).f\_magic** macro), the calling application can always determine whether the *FileHeader* pointer should refer to a 32-bit FILHDR or 64-bit FILHDR\_64 structure.

### Parameters

<i>ldPointer</i>	Points to the <b>LDFILE</b> structure that was returned as the result of a successful call to <b>ldopen</b> or <b>ldaopen</b> subroutine.
<i>FileHeader</i>	Points to a buffer large enough to accommodate a <b>FILHDR</b> structure, according to the object mode of the file being read.

### Return Values

The **ldfhread** subroutine returns Success or Failure.

### Error Codes

The **ldfhread** subroutine fails if it cannot read the file header.

**Note:** In most cases, the use of **ldfhread** can be avoided by using the **HEADER (ldPointer)** macro defined in the **ldfcn.h** file. The information in any field or fieldname of the header file may be accessed using the **header (ldPointer) fieldname** macro.

### Examples

The following is an example of code that opens an object file, determines its mode, and uses the **ldfhread** subroutine to acquire the file header. This code would be compiled with both **\_XCOFF32\_** and **\_XCOFF64\_** defined:

```
#define _XCOFF32_
#define __XCOFF64__

#include <ldfcn.h>
```

```

/* for each FileName to be processed */
if ( (ldPointer = ldopen(fileName, ldPointer)) != NULL)
{
    FILHDR    FileHead32;
    FILHDR_64 FileHead64;
    void      *FileHeader;

    if ( HEADER(ldPointer).f_magic == U802TOCMAGIC )
        FileHeader = &FileHead32;
    else if ( HEADER(ldPointer).f_magic == U803XTOCMAGIC )
        FileHeader = &FileHead64;
    else
        FileHeader = NULL;

    if ( FileHeader && (ldfhread( ldPointer, FileHeader ) == SUCCESS) )
    {
        /* ...successfully read header... */
        /* ...process according to magic number... */
    }
}

```

## Related Information

The **ldahread** (“ldahread Subroutine” on page 696) subroutine, **ldgetname** (“ldgetname Subroutine”) subroutine, **ldlread**, **ldlinit**, or **ldlitem** (“ldlread, ldlinit, or ldlitem Subroutine” on page 704) subroutine, **ldopen** (“ldopen or ldaopen Subroutine” on page 707) subroutine, **ldshread** or **ldnshread** (“ldshread or ldnshread Subroutine” on page 710) subroutine, **ldtbread** (“ldtbread Subroutine” on page 714) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldgetname Subroutine

### Purpose

Retrieves symbol name for common object file symbol table entry.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```

#include <stdio.h>
#include <ldfcn.h>

```

```

char *ldgetname ( ldPointer, Symbol)
LDFILE *ldPointer;
void *Symbol;

```

### Description

The **ldgetname** subroutine returns a pointer to the name associated with *Symbol* as a string. The string is in a static buffer local to the **ldgetname** subroutine that is overwritten by each call to the **ldgetname** subroutine and must therefore be copied by the caller if the name is to be saved.

The common object file format handles arbitrary length symbol names with the addition of a string table. The **ldgetname** subroutine returns the symbol name associated with a symbol table entry for an XCOFF-format object file.

The calling routine to provide a pointer to a buffer large enough to contain a symbol table entry for the associated object file. Since the **ldopen** subroutine provides magic number information (via the **HEADER(ldPointer).f\_magic** macro), the calling application can always determine whether the Symbol pointer should refer to a 32-bit SYMENT or 64-bit SYMENT\_64 structure.

The maximum length of a symbol name is **BUFSIZ**, defined in the **stdio.h** file.

## Parameters

<i>ldPointer</i>	Points to an <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> or <b>ldaopen</b> subroutine.
<i>Symbol</i>	Points to an initialized 32-bit or 64-bit <b>SYMENT</b> structure.

## Error Codes

The **ldgetname** subroutine returns a null value (defined in the **stdio.h** file) for a COFF-format object file if the name cannot be retrieved. This situation can occur if one of the following is true:

- The string table cannot be found.
- The string table appears invalid (for example, if an auxiliary entry is handed to the **ldgetname** subroutine wherein the name offset lies outside the boundaries of the string table).
- The name's offset into the string table is past the end of the string table.

Typically, the **ldgetname** subroutine is called immediately after a successful call to the **ldtbread** subroutine to retrieve the name associated with the symbol table entry filled by the **ldtbread** subroutine.

## Examples

The following is an example of code that determines the object file type before making a call to the **ldtbread** and **ldgetname** subroutines.

```
#define __XCOFF32__
#define __XCOFF64__

#include <ldfcn.h>

SYMENT Symbol32;
SYMENT_64 Symbol64;
void *Symbol;

if ( HEADER(ldPointer).f_magic == U802TOCMAGIC )
    Symbol = &Symbol32;
else if ( HEADER(ldPointer).f_magic == U64_TOCMAGIC )
    Symbol = &Symbol64;
else
    Symbol = NULL;

if ( Symbol )
    /* for each symbol in the symbol table */
    for ( symnum = 0 ; symnum < HEADER(ldPointer).f_nsyms ; symnum++ )
    {
        if ( ldtbread(ldPointer,symnum,Symbol) == SUCCESS )
        {
            char *name = ldgetname(ldPointer,Symbol)

            if ( name )
            {
```

```

        /* Got the name... */
        .
    }

    /* Increment symnum by the number of auxiliary entries */
    if ( HEADER(ldPointer).f_magic == U802TOCMAGIC )
        symnum += Symbol32.n_numaux;
    else if ( HEADER(ldPointer).f_magic == U64_TOCMAGIC )
        symnum += Symbol64.n_numaux;
    }
    else
    {
        /* Should have been a symbol...indicate the error */
        .
    }
}

```

## Related Information

The **ldahread** (“ldahread Subroutine” on page 696) subroutine, **ldfhread** (“ldfhread Subroutine” on page 701) subroutine, **ldlread**, **ldlinit**, or **ldlitem** (“ldlread, ldlinit, or ldlitem Subroutine”) subroutine, **ldshread** or **ldnshread** (“ldshread or ldnshread Subroutine” on page 710) subroutine, **ldtbread** (“ldtbread Subroutine” on page 714) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldlread, ldlinit, or ldlitem Subroutine

### Purpose

Manipulates line number entries of a common object file function.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```

#include <stdio.h>
#include <ldfcn.h>

```

```

int ldlread ( ldPointer, FunctionIndex, LineNumber, LineEntry)
LDFILE *ldPointer;
int FunctionIndex;
unsigned short LineNumber;
void *LineEntry;

```

```

int ldlinit (ldPointer, FunctionIndex)
LDFILE *ldPointer;
int FunctionIndex;

```

```

int ldlitem (ldPointer, LineNumber, LineEntry)
LDFILE *ldPointer;
unsigned short LineNumber;
void *LineEntry;

```



## Description

The **ldread** subroutine searches the line number entries of the XCOFF file currently associated with the *ldPointer* parameter. The **ldread** subroutine begins its search with the line number entry for the beginning of a function and confines its search to the line numbers associated with a single function. The function is identified by the *FunctionIndex* parameter, the index of its entry in the object file symbol table. The **ldread** subroutine reads the entry with the smallest line number equal to or greater than the *LineNumber* parameter into the memory beginning at the *LineEntry* parameter. It is the responsibility of the calling routine to provide a pointer to a buffer large enough to contain the line number entry for the associated object file type. Since the **ldopen** subroutine provides magic number information (via the **HEADER(*ldPointer*),f\_magic** macro), the calling application can always determine whether the *LineEntry* pointer should refer to a 32-bit LINENO or 64-bit LINENO\_64 structure.

The **ldlinit** and **ldlitem** subroutines together perform the same function as the **ldread** subroutine. After an initial call to the **ldread** or **ldlinit** subroutine, the **ldlitem** subroutine may be used to retrieve successive line number entries associated with a single function. The **ldlinit** subroutine simply locates the line number entries for the function identified by the *FunctionIndex* parameter. The **ldlitem** subroutine finds and reads the entry with the smallest line number equal to or greater than the *LineNumber* parameter into the memory beginning at the *LineEntry* parameter.

## Parameters

<i>ldPointer</i>	Points to the <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> , <b>lddopen</b> , or <b>ldaopen</b> subroutine.
<i>LineNumber</i>	Specifies the index of the first <i>LineNumber</i> parameter entry to be read.
<i>LineEntry</i>	Points to a buffer that will be filled in with a <b>LINENO</b> structure from the object file.
<i>FunctionIndex</i>	Points to the symbol table index of a function.

## Return Values

The **ldread**, **ldlinit**, and **ldlitem** subroutines return a SUCCESS or FAILURE value.

## Error Codes

The **ldread** subroutine fails if there are no line number entries in the object file, if the *FunctionIndex* parameter does not index a function entry in the symbol table, or if it finds no line number equal to or greater than the *LineNumber* parameter. The **ldlinit** subroutine fails if there are no line number entries in the object file or if the *FunctionIndex* parameter does not index a function entry in the symbol table. The **ldlitem** subroutine fails if it finds no line number equal to or greater than the *LineNumber* parameter.

## Related Information

The **ldahread** (“ldahread Subroutine” on page 696) subroutine, **ldfhread** (“ldfhread Subroutine” on page 701) subroutine, **ldgetname** (“ldgetname Subroutine” on page 702) subroutine, **ldshread** or **ldnshread** (“ldshread or ldnshread Subroutine” on page 710) subroutine, **ldtbread** (“ldtbread Subroutine” on page 714) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldlseek or ldnlseek Subroutine

### Purpose

Seeks to line number entries of a section of a common object file.

## Library

Object File Access Routine Library (**libld.a**)

## Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldlseek ( ldPointer, SectionIndex)
LDFILE *ldPointer;
unsigned short SectionIndex;
```

```
int ldnlseek (ldPointer, SectionName)
LDFILE *ldPointer;
char *SectionName;
```

## Description

The **ldlseek** subroutine seeks to the line number entries of the section specified by the *SectionIndex* parameter of the common object file currently associated with the *ldPointer* parameter. The first section has an index of 1.

The **ldnlseek** subroutine seeks to the line number entries of the section specified by the *SectionName* parameter.

Both subroutines determine the object mode of the associated file before seeking to the relocation entries of the indicated section.

## Parameters

<i>ldPointer</i>	Points to the <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> or <b>ldaopen</b> subroutine.
<i>SectionIndex</i>	Specifies the index of the section whose line number entries are to be sought to.
<i>SectionName</i>	Specifies the name of the section whose line number entries are to be sought to.

## Return Values

The **ldlseek** and **ldnlseek** subroutines return a SUCCESS or FAILURE value.

## Error Codes

The **ldlseek** subroutine fails if the *SectionIndex* parameter is greater than the number of sections in the object file. The **ldnlseek** subroutine fails if there is no section name corresponding with the *SectionName* parameter. Either function fails if the specified section has no line number entries or if it cannot seek to the specified line number entries.

## Related Information

The **ldohseek** (“ldohseek Subroutine” on page 707) subroutine, **ldrseek** or **ldnrseek** (“ldrseek or ldnrseek Subroutine” on page 709) subroutine, **ldsseek** or **ldnsseek** (“ldsseek or ldnsseek Subroutine” on page 712) subroutine, **ldtbseek** (“ldtbseek Subroutine” on page 715) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldohseek Subroutine

### Purpose

Seeks to the optional file header of a common object file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldohseek ( ldPointer)
LDFILE *ldPointer;
```

### Description

The **ldohseek** subroutine seeks to the optional auxiliary header of the common object file currently associated with the *ldPointer* parameter. The subroutine determines the object mode of the associated file before seeking to the end of its file header.

### Parameters

*ldPointer*                Points to the **LDFILE** structure that was returned as the result of a successful call to **ldopen** or **ldaopen** subroutine.

### Return Values

The **ldohseek** subroutine returns a SUCCESS or FAILURE value.

### Error Codes

The **ldohseek** subroutine fails if the object file has no optional header, if the file is not a 32-bit or 64-bit object file, or if it cannot seek to the optional header.

### Related Information

The **ldlseek** or **ldnlseek** (“ldlseek or ldnlseek Subroutine” on page 705) subroutine, **ldrseek** or **ldnrseek** (“ldrseek or ldnrseek Subroutine” on page 709) subroutine, **ldsseek** or **ldnsseek** (“ldsseek or ldnsseek Subroutine” on page 712) subroutine, **ldtbseek** (“ldtbseek Subroutine” on page 715) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldopen or ldaopen Subroutine

### Purpose

Opens an object or archive file for reading.

### Library

Object File Access Routine Library (**libld.a**)

## Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
LDFILE *ldopen( FileName, ldPointer)
char *FileName;
LDFILE *ldPointer;
```

```
LDFILE *ldaopen(FileName, ldPointer)
char *FileName;
LDFILE *ldPointer;
```

```
LDFILE *lddopen(FileDescriptor, type, ldPointer)
int FileDescriptor;
char *type;
LDFILE *ldPointer;
```

## Description

The **ldopen** and **ldclose** subroutines provide uniform access to both simple object files and object files that are members of archive files. Thus, an archive of object files can be processed as if it were a series of ordinary object files.

If the *ldPointer* is null, the **ldopen** subroutine opens the file named by the *FileName* parameter and allocates and initializes an **LDFILE** structure, and returns a pointer to the structure.

If the *ldPointer* parameter is not null and refers to an **LDFILE** for an archive, the structure is updated for reading the next archive member. In this case, and if the value of the **TYPE(ldPointer)** macro is the archive magic number **ARTYPE**.

The **ldopen** and **ldclose** subroutines are designed to work in concert. The **ldclose** subroutine returns failure only when the *ldPointer* refers to an archive containing additional members. Only then should the **ldopen** subroutine be called with a num-null *ldPointer* argument. In all other cases, in particular whenever a new *FileName* parameter is opened, the **ldopen** subroutine should be called with a null *ldPointer* argument.

If the value of the *ldPointer* parameter is not null, the **ldaopen** subroutine opens the *FileName* parameter again and allocates and initializes a new **LDFILE** structure, copying the **TYPE**, **OFFSET**, and **HEADER** fields from the *ldPointer* parameter. The **ldaopen** subroutine returns a pointer to the new **ldfile** structure. This new pointer is independent of the old pointer, *ldPointer*. The two pointers may be used concurrently to read separate parts of the object file. For example, one pointer may be used to step sequentially through the relocation information, while the other is used to read indexed symbol table entries.

The **lddopen** function accesses the previously opened file referenced by the *FileDescriptor* parameter. In all other respects, it functions the same as the **ldopen** subroutine.

For AIX 4.3.2 and above, the functions transparently open both 32-bit and 64-bit object files, as well as both small format and large format archive files. Once a file or archive is successfully opened, the calling application can examine the **HEADER(ldPointer).f\_magic** field to check the magic number of the file or archive member associated with *ldPointer*. (This is necessary due to an archive potentially containing members that are not object files.) The magic numbers **U802TOCMAGIC** and (for AIX 4.3.2 and above) **U803XTOCMAGIC** are defined in the **ldfcn.h** file. If the value of **TYPE(ldPointer)** is the archive magic

number **ARTYPE**, the flags field can be checked for the archive type. Large format archives will have the flag bit **AR\_TYPE\_BIG** set in **LDFLAGS(ldPointer)**. Large format archives are available on AIX 4.3 and later.

## Parameters

<i>FileName</i>	Specifies the file name of an object file or archive.
<i>ldPointer</i>	Points to an <b>LDFILE</b> structure.
<i>FileDescriptor</i>	Specifies a valid open file descriptor.
<i>type</i>	Points to a character string specifying the mode for the open file. The <b>fdopen</b> function is used to open the file.

## Error Codes

Both the **ldopen** and **ldlopen** subroutines open the file named by the *FileName* parameter for reading. Both functions return a null value if the *FileName* parameter cannot be opened, or if memory for the **LDFILE** structure cannot be allocated.

A successful open does not ensure that the given file is a common object file or an archived object file.

## Examples

The following is an example of code that uses the **ldopen** and **ldclose** subroutines:

```
/* for each FileName to be processed */

ldPointer = NULL;
do
    if((ldPointer = ldopen(FileName, ldPointer)) != NULL)
        /* check magic number */
        /* process the file */
        "
        "
    while(ldclose(ldPointer) == FAILURE );
```

## Related Information

The **ldclose** or **ldaclose** (“ldclose or ldaclose Subroutine” on page 697) subroutine, **fopen**, **fopen64**, **freopen**, **freopen64**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldrseek or ldnrseek Subroutine

### Purpose

Seeks to the relocation entries of a section of an XCOFF file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldrseek ( ldPointer, SectionIndex)
ldfile *ldPointer;
unsigned short SectionIndex;
```

```
int ldnrseek (ldPointer, SectionName)
ldfile *ldPointer;
char *SectionName;
```

## Description

The **ldrseek** subroutine seeks to the relocation entries of the section specified by the *SectionIndex* parameter of the common object file currently associated with the *ldPointer* parameter.

The **ldnrseek** subroutine seeks to the relocation entries of the section specified by the *SectionName* parameter.

For AIX 4.3.2 and above, both subroutines determine the object mode of the associated file before seeking to the relocation entries of the indicated section.

## Parameters

<i>ldPointer</i>	Points to an <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> , <b>lddopen</b> , or <b>ldaopen</b> subroutines.
<i>SectionIndex</i>	Specifies an index for the section whose relocation entries are to be sought.
<i>SectionName</i>	Specifies the name of the section whose relocation entries are to be sought.

## Return Values

The **ldrseek** and **ldnrseek** subroutines return a SUCCESS or FAILURE value.

## Error Codes

The **ldrseek** subroutine fails if the contents of the *SectionIndex* parameter are greater than the number of sections in the object file. The **ldnrseek** subroutine fails if there is no section name corresponding with the *SectionName* parameter. Either function fails if the specified section has no relocation entries or if it cannot seek to the specified relocation entries.

**Note:** The first section has an index of 1.

## Related Information

The **ldohseek** (“ldohseek Subroutine” on page 707) subroutine, **ldlseek** or **ldnlseek** (“ldlseek or ldnlseek Subroutine” on page 705) subroutine, **ldsseek** or **ldnsseek** (“ldsseek or ldnsseek Subroutine” on page 712) subroutine, **ldtbseek** (“ldtbseek Subroutine” on page 715) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldshread or ldnshread Subroutine

### Purpose

Reads a section header of an XCOFF file.

## Library

Object File Access Routine Library (**libld.a**)

## Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldshread ( ldPointer, SectionIndex, SectionHead)
LDFILE *ldPointer;
unsigned short SectionIndex;
void *SectionHead;
```

```
int ldnsread (ldPointer, SectionName, SectionHead)
LDFILE *ldPointer;
char *SectionName;
void *SectionHead;
```

## Description

The **ldshread** subroutine reads the section header specified by the *SectionIndex* parameter of the common object file currently associated with the *ldPointer* parameter into the area of memory beginning at the location specified by the *SectionHead* parameter.

The **ldnsread** subroutine reads the section header named by the *SectionName* argument into the area of memory beginning at the location specified by the *SectionHead* parameter. It is the responsibility of the calling routine to provide a pointer to a buffer large enough to contain the section header of the associated object file. Since the **ldopen** subroutine provides magic number information (via the **HEADER(ldPointer).f\_magic** macro), the calling application can always determine whether the *SectionHead* pointer should refer to a 32-bit **SCNHDR** or 64-bit **SCNHDR\_64** structure.

Only the first section header named by the *SectionName* argument is returned by the **ldshread** subroutine.

## Parameters

<i>ldPointer</i>	Points to an <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> , <b>lldopen</b> , or <b>ldaopen</b> subroutine.
<i>SectionIndex</i>	Specifies the index of the section header to be read. <b>Note:</b> The first section has an index of 1.
<i>SectionHead</i>	Points to a buffer large enough to accept either a 32-bit or a 64-bit <b>SCNHDR</b> structure, according to the object mode of the file being read.
<i>SectionName</i>	Specifies the name of the section header to be read.

## Return Values

The **ldshread** and **ldnsread** subroutines return a **SUCCESS** or **FAILURE** value.

## Error Codes

The **ldshread** subroutine fails if the *SectionIndex* parameter is greater than the number of sections in the object file. The **ldnsread** subroutine fails if there is no section with the name specified by the *SectionName* parameter. Either function fails if it cannot read the specified section header.

## Examples

The following is an example of code that opens an object file, determines its mode, and uses the **ldnshread** subroutine to acquire the .text section header. This code would be compiled with both **\_\_XCOFF32\_\_** and **\_\_XCOFF64\_\_** defined:

```
#define __XCOFF32__
#define __XCOFF64__

#include <ldfcn.h>

/* for each FileName to be processed */

if ( (ldPointer = ldopen(FileName, ldPointer)) != NULL )
{
    SCNHDR    SectionHead32;
    SCNHDR_64 SectionHead64;
    void      *SectionHeader;

    if ( HEADER(ldPointer).f_magic == U802TOCMAGIC )
        SectionHeader = &SectionHead32;
    else if ( HEADER(ldPointer).f_magic == U803XTOCMAGIC )
        SectionHeader = &SectionHead64;
    else
        SectionHeader = NULL;

    if ( SectionHeader && (ldnshread( ldPointer, ".text", SectionHeader ) == SUCCESS) )
    {
        /* ...successfully read header... */
        /* ...process according to magic number... */
    }
}
```

## Related Information

The **ldahread** (“ldahread Subroutine” on page 696) subroutine, **ldfhread** (“ldfhread Subroutine” on page 701) subroutine, **ldgetname** (“ldgetname Subroutine” on page 702) subroutine, **ldlread**, **ldlinit**, or **ldlitem** (“ldlread, ldlinit, or ldlitem Subroutine” on page 704) subroutine, **ldtbread** (“ldtbread Subroutine” on page 714) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldsseek or ldnsseek Subroutine

### Purpose

Seeks to an indexed or named section of a common object file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldsseek ( ldPointer, SectionIndex)
```



```
LDFILE *ldPointer;  
unsigned short SectionIndex;
```

```
int ldnseek (ldPointer, SectionName)  
LDFILE *ldPointer;  
char *SectionName;
```

## Description

The **ldsseek** subroutine seeks to the section specified by the *SectionIndex* parameter of the common object file currently associated with the *ldPointer* parameter. The subroutine determines the object mode of the associated file before seeking to the indicated section.

The **ldnsseek** subroutine seeks to the section specified by the *SectionName* parameter.

## Parameters

<i>ldPointer</i>	Points to the <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> or <b>ldaopen</b> subroutine.
<i>SectionIndex</i>	Specifies the index of the section whose line number entries are to be sought to.
<i>SectionName</i>	Specifies the name of the section whose line number entries are to be sought to.

## Return Values

The **ldsseek** and **ldnsseek** subroutines return a SUCCESS or FAILURE value.

## Error Codes

The **ldsseek** subroutine fails if the *SectionIndex* parameter is greater than the number of sections in the object file. The **ldnsseek** subroutine fails if there is no section name corresponding with the *SectionName* parameter. Either function fails if there is no section data for the specified section or if it cannot seek to the specified section.

**Note:** The first section has an index of 1.

## Related Information

The **ldlseek** or **ldnlseek** (“ldlseek or ldnlseek Subroutine” on page 705) subroutine, **ldohseek** (“ldohseek Subroutine” on page 707) subroutine, **ldrseek** or **ldnrseek** (“ldrseek or ldnrseek Subroutine” on page 709) subroutine, **ldtbseek** (“ldtbseek Subroutine” on page 715) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldtbindex Subroutine

### Purpose

Computes the index of a symbol table entry of a common object file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>  
#include <ldfcn.h>
```

```
long ldtbindex ( ldPointer)
LDFILE *ldPointer;
```

## Description

The **ldtbindex** subroutine returns the index of the symbol table entry at the current position of the common object file associated with the *ldPointer* parameter.

The index returned by the **ldtbindex** subroutine may be used in subsequent calls to the **ldtbread** subroutine. However, since the **ldtbindex** subroutine returns the index of the symbol table entry that begins at the current position of the object file, if the **ldtbindex** subroutine is called immediately after a particular symbol table entry has been read, it returns the index of the next entry.

## Parameters

*ldPointer*                Points to the **LDFILE** structure that was returned as a result of a successful call to the **ldopen** or **ldaopen** subroutine.

## Return Values

The **ldtbindex** subroutine returns the value **BADINDEX** upon failure. Otherwise a value greater than or equal to zero is returned.

## Error Codes

The **ldtbindex** subroutine fails if there are no symbols in the object file or if the object file is not positioned at the beginning of a symbol table entry.

**Note:** The first symbol in the symbol table has an index of 0.

## Related Information

The **ldtbread** (“ldtbread Subroutine”) subroutine, **ldtbseek** (“ldtbseek Subroutine” on page 715) subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldtbread Subroutine

### Purpose

Reads an indexed symbol table entry of a common object file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldtbread ( ldPointer, SymbolIndex, Symbol)
LDFILE *ldPointer;
long SymbolIndex;
void *Symbol;
```

## Description

The **ldtbread** subroutine reads the symbol table entry specified by the *SymbolIndex* parameter of the common object file currently associated with the *ldPointer* parameter into the area of memory beginning at the *Symbol* parameter. It is the responsibility of the calling routine to provide a pointer to a buffer large enough to contain the symbol table entry of the associated object file. Since the *ldopen* subroutine provides magic number information (via the **HEADER(*ldPointer*),f\_magic** macro), the calling application can always determine whether the *Symbol* pointer should refer to a 32-bit **SYMENT** or 64-bit **SYMENT\_64** structure.

## Parameters

<i>ldPointer</i>	Points to the <b>LDFILE</b> structure that was returned as the result of a successful call to the <b>ldopen</b> or <b>ldlopen</b> subroutine.
<i>SymbolIndex</i>	Specifies the index of the symbol table entry to be read.
<i>Symbol</i>	Points to a either a 32-bit or a 64-bit <b>SYMENT</b> structure.

## Return Values

The **ldtbread** subroutine returns a SUCCESS or FAILURE value.

## Error Codes

The **ldtbread** subroutine fails if the *SymbolIndex* parameter is greater than or equal to the number of symbols in the object file, or if it cannot read the specified symbol table entry.

**Note:** The first symbol in the symbol table has an index of 0.

## Related Information

The **ldahread** (“ldahread Subroutine” on page 696) subroutine, **ldfhread** (“ldfhread Subroutine” on page 701) subroutine, **ldgetname** (“ldgetname Subroutine” on page 702) subroutine, **ldlread**, **ldlinit**, or **ldlitem** (“ldlread, ldlinit, or ldlitem Subroutine” on page 704) subroutine, **ldshread** or **ldnshread** (“ldshread or ldnshread Subroutine” on page 710) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ldtbseek Subroutine

### Purpose

Seeks to the symbol table of a common object file.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
#include <stdio.h>
#include <ldfcn.h>
```

```
int ldtbseek ( ldPointer)
LDFILE *ldPointer;
```

## Description

The **ldtbseek** subroutine seeks to the symbol table of the common object file currently associated with the *ldPointer* parameter.

## Parameters

*ldPointer* Points to the **LDFILE** structure that was returned as the result of a successful call to the **ldopen** or **ldaopen** subroutine.

## Return Values

The **ldtbseek** subroutine returns a SUCCESS or FAILURE value.

## Error Codes

The **ldtbseek** subroutine fails if the symbol table has been stripped from the object file or if the subroutine cannot seek to the symbol table.

## Related Information

The **ldlseek** or **ldnlseek** (“ldlseek or ldnlseek Subroutine” on page 705) subroutine, **ldohseek** (“ldohseek Subroutine” on page 707) subroutine, **ldrseek** or **ldnrseek** (“ldrseek or ldnrseek Subroutine” on page 709) subroutine, **ldsseek** or **ldnsseek** (“ldsseek or ldnsseek Subroutine” on page 712) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## Igamma, Igammaf, or Igamma1 Subroutine

### Purpose

Computes the log gamma.

### Syntax

```
#include <math.h>

extern int signgam;

double lgamma (x)
double x;

float lgammaf (x)
float x;

long double lgamma1 (x)
long double x;
```

### Description

The sign of Gamma ( *x* ) is returned in the external integer **signgam**.

The **lgamma**, **lgammaf**, and **lgamma1** subroutines are not reentrant. A function that is not required to be reentrant is not required to be thread-safe.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **lgamma**, **lgammaf**, and **lgammal** subroutines return the logarithmic gamma of *x*.

If *x* is a non-positive integer, a pole error shall occur and **lgamma**, **lgammaf**, and **lgammal** will return **+HUGE\_VAL**, **+HUGE\_VALF**, and **+HUGE\_VALL**.

If the correct value would cause overflow, a range error shall occur and **lgamma**, **lgammaf**, and **lgammal** will return **±HUGE\_VAL**, **±HUGE\_VALF**, **±HUGE\_VALL**, respectively.

If *x* is NaN, a NaN is returned.

If *x* is 1 or 2, +0 is returned.

If *x* is  $\pm\text{Inf}$ , +Inf is returned.

## Related Information

“exp, expf, or expl Subroutine” on page 245, “feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## lineout Subroutine

### Purpose

Formats a print line.

### Library

None (provided by the print formatter)

### Syntax

```
#include <piostruct.h>
```

```
int lineout ( fileptr)
```

```
FILE *fileptr;
```

### Description

The **lineout** subroutine is invoked by the formatter driver only if the **setup** subroutine returns a non-null pointer. This subroutine is invoked for each line of the document being formatted. The **lineout** subroutine reads the input data stream from the *fileptr* parameter. It then formats and outputs the print line until it recognizes a situation that causes vertical movement on the page.

The **lineout** subroutine should process all characters to be printed and all printer commands related to horizontal movement on the page.

The **lineout** subroutine should not output any printer commands that cause vertical movement on the page. Instead, it should update the **vpos** (new vertical position) variable pointed to by the **shars\_vars**

structure that it shares with the formatter driver to indicate the new vertical position on the page. It should also refresh the **shar\_vars** variables for vertical increment and vertical decrement (reverse line-feed) commands.

When the **lineout** subroutine returns, the formatter driver sends the necessary commands to the printer to advance to the new vertical position on the page. This position is specified by the **vpos** variable. The formatter driver automatically handles top and bottom margins, new pages, initial pages to be skipped, and progress reports to the **qdaemon** daemon.

The following conditions can cause vertical movements:

- Line-feed control character or variable line-feed control sequence
- Vertical-tab control character
- Form-feed control character
- Reverse line-feed control character
- A line too long for the printer that wraps to the next line

Other conditions unique to a specific printer also cause vertical movement.

## Parameters

*fileptr* Specifies a file structure for the input data stream.

## Return Values

Upon successful completion, the **lineout** subroutine returns the number of bytes processed from the input data stream. It excludes the end-of-file character and any control characters or escape sequences that result only in vertical movement on the page (for example, line feed or vertical tab).

If a value of 0 is returned and the value in the **vpos** variable pointed to by the **shars\_vars** structure has not changed, or there are no more data bytes in the input data stream, the formatter driver assumes that printing is complete.

If the **lineout** subroutine detects an error, it uses the **piomsgout** subroutine to issue an error message. It then invokes the **pioexit** subroutine with a value of PIOEXITBAD.

**Note:** If either the **piocmdout** or **piogetstr** subroutine detects an error, it automatically issues its own error messages and terminates the print job.

## Related Information

The **piocmdout** subroutine, **pioexit** subroutine, **piogetstr** subroutine, **piomsgout** subroutine, **setup** subroutine.

Adding a New Printer Type to Your System and Printer Addition Management Subsystem: Programming Overview in *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*.

Print formatter example in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## link Subroutine

### Purpose

Creates an additional directory entry for an existing file.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <unistd.h>
```

```
int link ( Path1, Path2)  
const char *Path1, *Path2;
```

## Description

The **link** subroutine creates an additional hard link (directory entry) for an existing file. Both the old and the new links share equal access rights to the underlying object.

## Parameters

*Path1* Points to the path name of an existing file.  
*Path2* Points to the path name of the directory entry to be created.

### Notes:

1. If Network File System (NFS) is installed on your system, these paths can cross into another node.
2. With hard links, both the *Path1* and *Path2* parameters must reside on the same file system. If *Path1* is a symbolic link, an error is returned. Creating links to directories requires root user authority.

## Return Values

Upon successful completion, the **link** subroutine returns a value of 0. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **link** subroutine is unsuccessful if one of the following is true:

<b>EACCES</b>	Indicates the requested link requires writing in a directory that denies write permission.
<b>EDQUOT</b>	Indicates the directory in which the entry for the new link is being placed cannot be extended, or disk blocks could not be allocated for the link because the user or group quota of disk blocks or i-nodes on the file system containing the directory has been exhausted.
<b>EEXIST</b>	Indicates the link named by the <i>Path2</i> parameter already exists.
<b>EMLINK</b>	Indicates the file already has the maximum number of links.
<b>ENOENT</b>	Indicates the file named by the <i>Path1</i> parameter does not exist.
<b>ENOSPC</b>	Indicates the directory in which the entry for the new link is being placed cannot be extended because there is no space left on the file system containing the directory.
<b>EPERM</b>	Indicates the file named by the <i>Path1</i> parameter is a directory, and the calling process does not have root user authority.
<b>EROFS</b>	Indicates the requested link requires writing in a directory on a read-only file system.
<b>EXDEV</b>	Indicates the link named by the <i>Path2</i> parameter and the file named by the <i>Path1</i> parameter are on different file systems, or the file named by <i>Path1</i> refers to a named STREAM.

The **link** subroutine can be unsuccessful for other reasons. See Appendix A, “Base Operating System Error Codes for Services That Require Path-Name Resolution,” on page 1345 for a list of additional errors.

If NFS is installed on the system, the **link** subroutine is unsuccessful if the following is true:

**ETIMEDOUT** Indicates the connection timed out.

## Related Information

The **symlink** subroutine, **unlink** subroutine.

The **link** or **unlink** command, **ln** command, **rm** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## lio\_listio or lio\_listio64 Subroutine

The **lio\_listio** or **lio\_listio64** subroutine includes information for the POSIX AIO **lio\_listio** subroutine (as defined in the IEEE std 1003.1-2001), and the Legacy AIO **lio\_listio** subroutine.

### POSIX AIO lio\_listio Subroutine

## Purpose

Initiates a list of asynchronous I/O requests with a single call.

## Syntax

```
#include <aio.h>
int lio_listio(mode, list, nent, sig)
int mode;
struct aiocb *restrict const list[restrict];
int nent;
struct sigevent *restrict sig;
```

## Description

The *lio\_listio* subroutine initiates a list of I/O requests with a single function call.

The *mode* parameter takes one of the values (LIO\_WAIT, LIO\_NOWAIT or LIO\_NOWAIT\_AIOWAIT) declared in **<aio.h>** and determines whether the subroutine returns when the I/O operations have been completed, or as soon as the operations have been queued. If the *mode* parameter is set to LIO\_WAIT, the subroutine waits until all I/O is complete and the *sig* parameter is ignored.

If the *mode* parameter is set to LIO\_NOWAIT or LIO\_NOWAIT\_AIOWAIT, the subroutine returns immediately. If LIO\_NOWAIT is set, asynchronous notification occurs, according to the *sig* parameter, when all I/O operations complete. If *sig* is NULL, no asynchronous notification occurs. If *sig* is not NULL, asynchronous notification occurs when all the requests in *list* have completed. If LIO\_NOWAIT\_AIOWAIT is set, the **aiocb\_nwait** subroutine must be called for the aio control blocks to be updated. For more information, see the “aiocb\_nwait Subroutine” on page 47.

The I/O requests enumerated by *list* are submitted in an unspecified order.

The *list* parameter is an array of pointers to **aiocb** structures. The array contains *nent* elements. The array may contain NULL elements, which are ignored.

The *aio\_lio\_opcode* field of each **aiocb** structure specifies the operation to be performed. The supported operations are LIO\_READ, LIO\_WRITE, and LIO\_NOP; these symbols are defined in **<aio.h>**. The LIO\_NOP operation causes the list entry to be ignored. If the *aio\_lio\_opcode* element is equal to LIO\_READ, an I/O operation is submitted as if by a call to **aio\_read** with the *aio\_cbp* equal to the address of the **aiocb** structure. If the *aio\_lio\_opcode* element is equal to LIO\_WRITE, an I/O operation is submitted as if by a call to **aio\_write** with the *aio\_cbp* argument equal to the address of the **aiocb** structure.

The *aio\_fildes* member specifies the file descriptor on which the operation is to be performed.



The *aio\_buf* member specifies the address of the buffer to or from which the data is transferred.

The *aio\_nbytes* member specifies the number of bytes of data to be transferred.

The members of the **aio\_cb** structure further describe the I/O operation to be performed, in a manner identical to that of the corresponding **aio\_cb** structure when used by the **aio\_read** and **aio\_write** subroutines.

The *nent* parameter specifies how many elements are members of the list.

The behavior of the **lio\_listio** subroutine is altered according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion if synchronized I/O is enabled on the file associated with *aio\_fildes* .

For regular files, no data transfer occurs past the offset maximum established in the open file description.

## Parameters

<i>mode</i>	Determines whether the subroutine returns when the I/O operations are completed, or as soon as the operations are queued.
<i>list</i>	An array of pointers to aio control structures defined in the <b>aio.h</b> file.
<i>nent</i>	Specifies the length of the array.
<i>sig</i>	Determines when asynchronous notification occurs.

## Execution Environment

The **lio\_listio** and **lio\_listio64** subroutines can be called from the process environment only.

## Return Values

<b>EAGAIN</b>	The resources necessary to queue all the I/O requests were not available. The application may check the error status of each <b>aio_cb</b> to determine the individual request(s) that failed.  The number of entries indicated by <i>nent</i> would cause the system-wide limit (AIO_MAX) to be exceeded.
<b>EINVAL</b>	The <i>mode</i> parameter is not a proper value, or the value of <i>nent</i> was greater than AIO_LISTIO_MAX.
<b>EINTR</b>	A signal was delivered while waiting for all I/O requests to complete during an LIO_WAIT operation. Since each I/O operation invoked by the <b>lio_listio</b> subroutine may provoke a signal when it completes, this error return may be caused by the completion of one (or more) of the very I/O operations being awaited. Outstanding I/O requests are not canceled, and the application examines each list element to determine whether the request was initiated, canceled, or completed.
<b>EIO</b>	One or more of the individual I/O operations failed. The application may check the error status for each <b>aio_cb</b> structure to determine the individual request(s) that failed.

If the **lio\_listio** subroutine succeeds or fails with errors of **EAGAIN**, **EINTR**, or **EIO**, some of the I/O specified by the list may have been initiated. If the **lio\_listio** subroutine fails with an error code other than **EAGAIN**, **EINTR**, or **EIO**, no operations from the list were initiated. The I/O operation indicated by each list element can encounter errors specific to the individual read or write function being performed. In this event, the error status for each *aio\_cb* control block contains the associated error code. The error codes that can be set are the same as would be set by the **read** or **write** subroutines, with the following additional error codes possible:

<b>EAGAIN</b>	The requested I/O operation was not queued due to resource limitations.
---------------	---

<b>ECANCELED</b>	The requested I/O was canceled before the I/O completed due to an <b>aio_cancel</b> request.
<b>EFBIG</b>	The <i>aio_lio_opcode</i> argument is LIO_WRITE, the file is a regular file, <i>aio_nbytes</i> is greater than 0, and <i>aio_offset</i> is greater than or equal to the offset maximum in the open file description associated with <i>aio_fildes</i> .
<b>EINPROGRESS</b>	The requested I/O is in progress.
<b>EOVERFLOW</b>	The <i>aio_lio_opcode</i> argument is set to LIO_READ, the file is a regular file, <i>aio_nbytes</i> is greater than 0, and the <i>aio_offset</i> argument is before the end-of-file and is greater than or equal to the offset maximum in the open file description associated with <i>aio_fildes</i> .

## Related Information

“aio\_cancel or aio\_cancel64 Subroutine” on page 38, “aio\_error or aio\_error64 Subroutine” on page 42, “aio\_read or aio\_read64 Subroutine” on page 50, “aio\_return or aio\_return64 Subroutine” on page 55, “aio\_suspend or aio\_suspend64 Subroutine” on page 58, “aio\_write or aio\_write64 Subroutine” on page 61, “close Subroutine” on page 175, “exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236, “exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243, “fork, f\_fork, or vfork Subroutine” on page 291, and “lseek, llseek or lseek64 Subroutine” on page 765.

The read, readx, readv, readvx, or pread Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

### Legacy AIO lio\_listio Subroutine

## Purpose

Initiates a list of asynchronous I/O requests with a single call.

## Syntax

```
#include <aio.h>
```

```
int lio_listio (cmd,
list, nent, eventp)
int cmd, nent;
struct liocb * list[ ];
struct event * eventp;

int lio_listio64
(cmd, list, nent, eventp)
int cmd, nent; struct liocb64 *list;
struct event *eventp;
```

## Description

The **lio\_listio** subroutine allows the calling process to initiate the *nent* parameter asynchronous I/O requests. These requests are specified in the **liocb** structures pointed to by the elements of the *list* array. The call may block or return immediately depending on the *cmd* parameter. If the *cmd* parameter requests that I/O completion be asynchronously notified, a **SIGIO** signal is delivered when all I/O operations are completed.

The **lio\_listio64** subroutine is similar to the **lio\_listio** subroutine except that it takes an array of pointers to **liocb64** structures. This allows the **lio\_listio64** subroutine to specify offsets in excess of OFF\_MAX (2 gigabytes minus 1).

In the large file enabled programming environment, **lio\_listio** is redefined to be **lio\_listio64**.

**Note:** The pointer to the **event** structure *eventp* parameter is currently not in use, but is included for future compatibility.

## Parameters

*cmd* The *cmd* parameter takes one of the following values:

### **LIO\_WAIT**

Queues the requests and waits until they are complete before returning.

### **LIO\_NOWAIT**

Queues the requests and returns immediately, without waiting for them to complete. The *event* parameter is ignored.

### **LIO\_NOWAIT\_AIOWAIT**

Queues the requests and returns immediately, without waiting for them to complete. The **aio\_nwait** subroutine must be called for the aio control blocks to be updated.

### **LIO\_ASYNC**

Queues the requests and returns immediately, without waiting for them to complete. An enhanced signal is delivered when all the operations are completed. Currently this command is not implemented.

### **LIO\_ASIG**

Queues the requests and returns immediately, without waiting for them to complete. A **SIGIO** signal is generated when all the I/O operations are completed.

*list* Points to an array of pointers to **liocb** structures. The structure array contains *nent* elements:

#### *lio\_aiocb*

The asynchronous I/O control block associated with this I/O request. This is an actual **aioacb** structure, not a pointer to one.

#### *lio\_fildes*

Identifies the file object on which the I/O is to be performed.

#### *lio\_opcode*

This field may have one of the following values defined in the **/usr/include/sys/aio.h** file:

### **LIO\_READ**

Indicates that the read I/O operation is requested.

### **LIO\_WRITE**

Indicates that the write I/O operation is requested.

### **LIO\_NOP**

Specifies that no I/O is requested (that is, this element will be ignored).

*nent* Specifies the number of entries in the array of pointers to **listio** structures.

*eventp* Points to an **event** structure to be used when the *cmd* parameter is set to the **LIO\_ASYNC** value. This parameter is currently ignored.

## Execution Environment

The **lio\_listio** and **lio\_listio64** subroutines can be called from the process environment only.

## Return Values

When the **lio\_listio** subroutine is successful, it returns a value of 0. Otherwise, it returns a value of -1 and sets the **errno** global variable to identify the error. The returned value indicates the success or failure of the **lio\_listio** subroutine itself and not of the asynchronous I/O requests (except when the command is **LIO\_WAIT**). The **aio\_error** subroutine returns the status of each I/O request.

If the **lio\_listio** subroutine succeeds or fails with errors of **EAGAIN**, **EINTR**, or **EIO**, some of the I/O specified by the *list* might have been initiated. If the **lio\_listio** subroutine fails with an error code other than **EAGAIN**, **EINTR**, or **EIO**, no operations from the *list* were initiated. The I/O operation indicated by each *list* element can encounter errors specific to the individual read or write function being performed. In this

event, the error status for each **aio** control block contains the associated error code. The error codes that can be set are the same as would be set by the read or write subroutines, with the following additional error codes possible:

<b>EAGAIN</b>	Indicates that the system resources required to queue the request are not available. Specifically, the transmit queue may be full, or the maximum number of opens may have been reached.
<b>EINTR</b>	Indicates that a signal or event interrupted the <b>lio_listio</b> subroutine call.
<b>EINVAL</b>	Indicates that the <code>aio_whence</code> field does not have a valid value or that the resulting pointer is not valid.
<b>EIO</b>	One or more of the individual I/O operations failed. The application can check the error status for each <b>aio</b> structure to determine the individual request that failed.

## Related Information

The **aio\_cancel** or **aio\_cancel64** (“aio\_cancel or aio\_cancel64 Subroutine” on page 38) subroutine, **aio\_error** or **aio\_error64** (“aio\_error or aio\_error64 Subroutine” on page 42) subroutine, **aio\_read** or **aio\_read64** (“aio\_read or aio\_read64 Subroutine” on page 50) subroutine, **aio\_return** or **aio\_return64** (“aio\_return or aio\_return64 Subroutine” on page 55) subroutine, **aio\_suspend** or **aio\_suspend64** (“aio\_suspend or aio\_suspend64 Subroutine” on page 58) subroutine, **aio\_write** or **aio\_write64** (“aio\_write or aio\_write64 Subroutine” on page 61) subroutine.

The Asynchronous I/O Overview and the Communications I/O Subsystem: Programming Introduction in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* describes the files, commands, and subroutines used for low-level, stream, terminal, and asynchronous I/O interfaces.

---

## listea Subroutine

### Purpose

Lists the extended attributes associated with a file.

### Syntax

```
#include <sys/ea.h>
```

```
ssize_t listea(const char *path, char *list, size_t size);
ssize_t flistea(int filedes, char *list, size_t size);
ssize_t llistea(const char *path, char *list, size_t size);
```

### Description

Extended attributes are name:value pairs associated with the file system objects (such as files, directories, and symlinks). They are extensions to the normal attributes that are associated with all objects in the file system (that is, the **stat(2)** data).

Do not define an extended attribute name with eight characters prefix "(0xF8)SYSTEM(0xF8)". Prefix "(0xF8)SYSTEM(0xF8)" is reserved for system use only.

**Note:** The 0xF8 prefix represents a non-printable character.

The **listea** subroutine retrieves the list of extended attribute names associated with the given *path* in the file system. The *list* is the set of (NULL-terminated) names, one after the other. Names of extended attributes to which the calling process does not have access might be omitted from the list. The length of the attribute name list is returned. The **flistea** subroutine is identical to **listea**, except that it takes a file descriptor instead of a path. The **llistea** subroutine is identical to **listea**, except, in the case of a symbolic link, the link itself is interrogated, not the file that it refers to.

An empty buffer of size 0 can be passed into these calls to return the current size of the list of extended attribute names, which can be used to estimate whether the size of a buffer is sufficiently large to hold the list of names.

## Parameters

<i>path</i>	The path name of the file.
<i>list</i>	A pointer to a buffer in which the list of attributes will be stored.
<i>size</i>	The size of the buffer.
<i>filedes</i>	A file descriptor for the file.

## Return Values

If the **listea** subroutine succeeds, a nonnegative number is returned that indicates the length in bytes of the attribute name list. Upon failure, -1 is returned and **errno** is set appropriately.

## Error Codes

<b>EACCES</b>	Caller lacks read permission on the base file, or lacks the appropriate ACL privileges for named attribute <b>read</b> .
<b>EFAULT</b>	A bad address was passed for <i>path</i> or <i>list</i> .
<b>EFORMAT</b>	File system is capable of supporting EAs, but EAs are disabled.
<b>ENOTSUP</b>	Extended attributes are not supported by the file system.
<b>ERANGE</b>	The size of the value buffer is too small to hold the result.

## Related Information

“getea Subroutine” on page 363, removeea Subroutine, setea Subroutine, stateea Subroutine.

---

## llrint, llrintf, llrintl, llrintd32, llrintd64, and llrintd128 Subroutines

### Purpose

Round to the nearest integer value using current rounding direction.

### Syntax

```
#include <math.h>
```

```
long long llrint (x)  
double x;
```

```
long long llrintf (x)  
float x;
```

```
long long llrintl (x)  
long double x;
```

```
long long llrintd32(x)  
_Decimal32 x;
```

```
long long llrintd64(x)  
_Decimal64 x;
```

```
long long llrintd128(x)  
_Decimal128 x;
```

## Description

The **llrint**, **llrintf**, **llrintl**, **llrintd32**, **llrintd64**, and **llrintd128** subroutines round the *x* parameter to the nearest integer value, according to the current rounding direction.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be rounded.

## Return Values

Upon successful completion, the **llrint**, **llrintf**, **llrintl**, **llrintd32**, **llrintd64**, and **llrintd128** subroutines return the rounded integer value.

If *x* is NaN, a domain error occurs, and an unspecified value is returned.

If *x* is +Inf, a domain error occurs and an unspecified value is returned.

If *x* is -Inf, a domain error occurs and an unspecified value is returned.

If the correct value is positive and too large to represent as a **long long**, a domain error occur and an unspecified value is returned.

If the correct value is negative and too large to represent as a **long long**, a domain error occurs and an unspecified value is returned.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## llround, llroundf, llroundl, llroundd32, llroundd64, and llroundd128 Subroutines

### Purpose

Round to the nearest integer value.

### Syntax

```
#include <math.h>
```

```
long long llround (x)  
double x;
```

```
long long llroundf (x)  
float x;
```

```
long long llroundl (x)  
long double x;
```

```
long long llroundd32(x)
```

```
_Decimal132 x;
```

```
long long llroundd64(x)
```

```
_Decimal64 x;
```

```
long long llroundd128(x)
```

```
_Decimal128 x;
```

## Description

The **llround**, **llroundf**, **llroundl**, **llroundd32**, **llroundd64**, and **llroundd128** subroutines round the *x* parameter to the nearest integer value, rounding halfway cases away from zero, regardless of the current rounding direction.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be rounded.

## Return Values

Upon successful completion, the **llround**, **llroundf**, **llroundl**, **llroundd32**, **llroundd64**, and **llroundd128** subroutines return the rounded integer value.

If *x* is NaN, a domain error occurs, and an unspecified value is returned.

If *x* is +Inf, a domain error occurs and an unspecified value is returned.

If *x* is -Inf, a domain error occurs and an unspecified value is returned.

If the correct value is positive and too large to represent as a **long long**, a domain error occurs and an unspecified value is returned.

If the correct value is negative and too large to represent as a **long long**, a domain error occurs and an unspecified value is returned.

## Related Information

“**feclearexcept** Subroutine” on page 265 and “**fetestexcept** Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## load and loadAndInit Subroutines

### Purpose

Loads a module into the current process.

### Syntax

```
int *load ( ModuleName, Flags, LibraryPath)
char *ModuleName;
uint Flags;
char *LibraryPath;
```

```
int *loadAndInit ( ModuleName, Flags, LibraryPath)
char *ModuleName;
uint Flags;
char *LibraryPath;
```

## Description

The **load** and **loadAndInit** subroutines load the specified module into the calling process's address space. A module can be a regular file or a member of an archive. When adding a new module to the address space of a 32-bit process, the load operation may cause the break value to change.

The **load** subroutine is not a preferred method to load C++ modules. Use **loadAndInit** subroutine instead. The **loadAndInit** subroutine uses the same interface as **load** but performs C++ initialization.

The **exec** subroutine is similar to the **load** subroutine, except that:

- The **load** subroutine does not replace the current program with a new one.
- The **exec** subroutine does not have an explicit library path parameter; it has only the **LIBPATH** and **LD\_LIBRARY\_PATH** environment variables. Also, these library path environment variables are ignored when the program using the **exec** subroutine has more privilege than the caller (for example, in the case of a **set-UID** program).

A large application can be split up into one or more modules in one of two ways that allow execution within the same process. The first way is to create each of the application's modules separately and use **load** to explicitly load a module when it is needed. The other way is to specify the relationship between the modules when they are created by defining imported and exported symbols.

Modules can import symbols from other modules. Whenever symbols are imported from one or more other modules, these modules are automatically loaded to resolve the symbol references if the required modules are not already loaded, and if the imported symbols are not specified as deferred imports. These modules can be archive members in libraries or individual files and can have either shared or private file characteristics that control how and where they are loaded.

Shared modules (typically members of a shared library archive) are loaded into the shared library region, when their access permissions allow sharing, that is, when they have read-other permission. Private modules, and shared modules without the required permissions for sharing, are loaded into the process private region.

When the loader resolves a symbol, it uses the file name recorded with that symbol to find the module that exports the symbol. If the file name contains any / (slash) characters, it is used directly and must name an appropriate file or archive member. However, if the file name is a base name (contains no / characters), the loader searches the directories specified in the default library path for a file (i.e. a module or an archive) with that base name.

The *LibraryPath* is a string containing one or more directory path names separated by colons. See the section "Searching for Dependent Modules" on page 729 for information on library path searching.

(This paragraph only applies to AIX 4.3.1 and previous releases.) When a process is executing under **ptrace** control, portions of the process's address space are recopied after the **load** processing completes. For a 32-bit process, the main program text (loaded in segment 1) and shared library modules (loaded in segment 13) are recopied. Any breakpoints or other modifications to these segments must be reinserted after the **load** call. For a 64-bit process, shared library modules are recopied after a **load** call. The debugger will be notified by setting the **W\_SLWTED** flag in the status returned by **wait**, so that it can reinsert breakpoints.

(This paragraph only applies to AIX 4.3.2 and later releases.) When a process executing under **ptrace** control calls **load**, the debugger is notified by setting the **W\_SLWTED** flag in the status returned by **wait**.



Any modules newly loaded into the shared library segments will be copied to the process's private copy of these segments, so that they can be examined or modified by the debugger.

If the program calling the **load** subroutine was linked on AIX 4.2 or a later release, the **load** subroutine will call initialization routines (**init** routines) for the new module and any of its dependents if they were not already loaded.

Modules loaded by this subroutine are automatically unloaded when the process terminates or when the **exec** subroutine is executed. They are explicitly unloaded by calling the **unload** subroutine.

## Searching for Dependent Modules

The load operation and the exec operation differ slightly in their dependent module search mechanism. When a module is added to the address space of a running process (the load operation), the rules outlined in the next section are used to find the named module. Note that dependency relationships may be loosely defined as a tree but recursive relationships between modules may also exist. The following components may be used to create a complete library search path:

1. If the **L\_LIBPATH\_EXEC** flag is set, the library search path used at exec-time.
2. The value of the *LibraryPath* parameter if it is non-null. Note that a null string is a valid search path which refers to the current working directory. If the *LibraryPath* parameter is NULL, the value of the **LIBPATH** environment variable, or alternatively the **LD\_LIBRARY\_PATH** environment variable (if **LIBPATH** is not set), is used instead.
3. The library search path contained in the loader section of the module being loaded (the *ModuleName* parameter).
4. The library search path contained in the loader section of the module whose immediate dependents are being loaded. Note that this per-module information changes when searching for each module's immediate dependents.

To find the *ModuleName* module, components 1 and 2 are used. To find dependents, components 1, 2, 3 and 4 are used in order. Note that if any modules that are already part of the running process satisfy the dependency requirements of the newly loaded module(s), pre-existing modules are not loaded again.

For each colon-separated portion of the aggregate search specification, if the base name is not found the search continues. The first instance of the base name found is used; if the file is not of the proper form, or in the case of an archive does not contain the required archive member, or does not export a definition of a required symbol, an error occurs. The library path search is not performed when either a relative or an absolute path name is specified for a dependent module.

The library search path stored within the module is specified at link-edit time.

The **load** subroutine may cause the calling process to fail if the module specified has a very long chain of dependencies, (for example, lib1.a, which depends on lib2.a, which depends on lib3.a, etc). This is because the loader processes such relationships recursively on a fixed-size stack. This limitation is exposed only when processing a dependency chain that has over one thousand elements.

## Parameters

<i>ModuleName</i>	<p>Points to the name of the module to be loaded. The module name consists of a path name, and, an optional member name. If the path name contains at least one / character, the name is used directly, and no directory searches are performed to locate the file. If the path name contains no / characters, it is treated as a base name, and should be in one of the directories listed in the library path.</p> <p>The library path is either the value of the <i>LibraryPath</i> parameter if not a null value, or the value of the <b>LIBPATH</b> environment variable (if set; otherwise, <b>LD_LIBRARY_PATH</b> environment variable, if set) or the library path used at process exec time (if the <b>L_LIBPATH_EXEC</b> is set). If no library path is provided, the module should be in the current directory.</p> <p>The <i>ModuleName</i> parameter may explicitly name an archive member. The syntax is <i>pathname(member)</i> where <i>pathname</i> follows the rules specified in the previous paragraph, and <i>member</i> is the name of a specific archive member. The parentheses are a required portion of the specification and no intervening spaces are allowed. If an archive member is named, the <b>L_LOADMEMBER</b> flag must be added to the <i>Flags</i> parameter. Otherwise, the entire <i>ModuleName</i> parameter is treated as an explicit filename.</p>
<i>Flags</i>	<p>Modifies the behavior of the <b>load</b> and the <b>loadAndInit</b> services as follows (see the <b>ldr.h</b> file). If no special behavior is required, set the value of the flags parameter to 0 (zero). For compatibility, a value of 1 (one) may also be specified.</p> <p><b>L_LIBPATH_EXEC</b></p> <p>Specifies that the library path used at process exec time should be prepended to any library path specified in the <b>load</b> call (either as an argument or environment variable). It is recommended that this flag be specified in all calls to the <b>load</b> subroutine.</p> <p><b>L_LOADMEMBER</b></p> <p>Indicates that the <i>ModuleName</i> parameter may specify an archive member. The <i>ModuleName</i> argument is searched for parentheses, and if found the parameter is treated as a filename/member name pair. If this flag is present and the <i>ModuleName</i> parameter does not contain parenthesis the entire <i>ModuleName</i> parameter is treated as a filename specification. Under either condition the filename is expected to be found within the library path or the current directory.</p> <p><b>L_NOAUTODEFER</b></p> <p>Specifies that any deferred imports in the module being loaded must be explicitly resolved by use of the <b>loadbind</b> subroutine. This allows unresolved imports to be explicitly resolved at a later time with a specified module. If this flag is not specified, deferred imports (marked for deferred resolution) are resolved at the earliest opportunity when any subsequently loaded module exports symbols matching unresolved imports.</p>
<i>LibraryPath</i>	<p>Points to a character string that specifies the default library search path.</p> <p>If the <i>LibraryPath</i> parameter is NULL, the <b>LIBPATH</b> environment variable is used, if set; otherwise, the <b>LD_LIBRARY_PATH</b> environment variable is used. See the section “Searching for Dependent Modules” on page 729 for more information.</p> <p>The library path is used to locate dependent modules that are specified as basenames (that is, their pathname components do not contain a / (slash) character).</p> <p>Note the difference between setting the <i>LibraryPath</i> parameter to null, and having the <i>LibraryPath</i> parameter point to a null string (“”). A null string is a valid library path which consists of a single directory: the current directory.</p>

## Return Values

Upon successful completion, the **load** and **loadAndInit** subroutines return the pointer to function for the entry point of the module. If the module has no entry point, the address of the data section of the module is returned.

## Error Codes

If the **load** and **loadAndInit** subroutines fail, a null pointer is returned, the module is not loaded, and **errno** global variable is set to indicate the error. The **load** and **loadAndInit** subroutines fail if one or more of the following are true of a module to be explicitly or automatically loaded:

<b>EACCES</b>	Indicates the file is not an ordinary file, or the mode of the program file denies execution permission, or search permission is denied on a component of the path prefix.
<b>EINVAL</b>	Indicates the file or archive member has a valid magic number in its header, but the header is damaged or is incorrect for the machine on which the file is to be run.
<b>ELOOP</b>	Indicates too many symbolic links were encountered in translating the path name.
<b>ENOEXEC</b>	Indicates an error occurred when loading or resolving symbols for the specified module. This can be due to an attempt to load a module with an invalid <b>XCOFF</b> header, a failure to resolve symbols that were not defined as deferred imports or several other load time related problems. The <b>loadquery</b> subroutine can be used to return more information about the load failure. If the main program was linked on a AIX 4.2 or later system, and if runtime linking is used, the <b>load</b> and the <b>loadAndInit</b> subroutines will fail if the runtime linker could not resolve some symbols. In this case, <b>errno</b> will be set to <b>ENOEXEC</b> , but the <b>loadquery</b> subroutine will not return any additional information.
<b>ENOMEM</b>	Indicates the program requires more memory than is allowed by the system-imposed maximum.
<b>ETXTBSY</b>	Indicates the file is currently open for writing by some process.
<b>ENAMETOOLONG</b>	Indicates a component of a path name exceeded 255 characters, or an entire path name exceeded 1023 characters.
<b>ENOENT</b>	Indicates a component of the path prefix does not exist, or the path name is a null value. For the <b>dlopen</b> subroutine, <b>RTLDMEMBER</b> is not used when trying to open a member within the archive file.
<b>ENOTDIR</b>	Indicates a component of the path prefix is not a directory.
<b>ESTALE</b>	Indicates the process root or current directory is located in a virtual file system that has been unmounted.

## Related Information

The **dlopen** (“dlopen Subroutine” on page 217) subroutine, **exec** (“exec: execl, execl, execlp, execlp, execlp, execlp, or exect Subroutine” on page 236) subroutine, **loadbind** (“loadbind Subroutine”) subroutine, **loadquery** (“loadquery Subroutine” on page 733) subroutine, **ptrace** (“ptrace, ptracex, ptracex64 Subroutine” on page 1307) subroutine, **unload** subroutine.

The **ld** command.

The Shared Library Overview and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

The Dynamically loading a shared library section in the *XL C/C++ V8.0 for AIX Programming Guide* book.

---

## loadbind Subroutine

### Purpose

Provides specific run-time resolution of a module’s deferred symbols.

### Syntax

```
int loadbind( Flag, ExportPointer, ImportPointer)
int Flag;
void *ExportPointer, *ImportPointer;
```

## Description

The **loadbind** subroutine controls the run-time resolution of a previously loaded object module's unresolved imported symbols.

The **loadbind** subroutine is used when two modules are loaded. Module A, an object module loaded at run time with the **load** subroutine, has designated that some of its imported symbols be resolved at a later time. Module B contains exported symbols to resolve module A's unresolved imports.

To keep module A's imported symbols from being resolved until the **loadbind** service is called, you can specify the **load** subroutine flag, **L\_NOAUTODEFER**, when loading module A.

(This paragraph only applies to AIX 4.3.1 and previous releases.) When a 32-bit process is executing under **ptrace** control, portions of the process's address space are recopied after the **loadbind** processing completes. The main program text (loaded in segment 1) and shared library modules (loaded in segment 13) are recopied. Any breakpoints or other modifications to these segments must be reinserted after the **loadbind** call.

(This paragraph only applies to AIX 4.3.2 and later releases.) When a 32-bit process executing under **ptrace** control calls **loadbind**, the debugger is notified by setting the **W\_SLWTED** flag in the status returned by **wait**.

When a 64-bit process under **ptrace** control calls **loadbind**, the debugger is not notified and execution of the process being debugged continues normally.

## Parameters

<i>Flag</i>	Currently not used.
<i>ExportPointer</i>	Specifies the function pointer returned by the <b>load</b> subroutine when module B was loaded.
<i>ImportPointer</i>	Specifies the function pointer returned by the <b>load</b> subroutine when module A was loaded.

**Note:** The *ImportPointer* or *ExportPointer* parameter may also be set to any exported static data area symbol or function pointer contained in the associated module. This would typically be the function pointer returned from the **load** of the specified module.

## Return Values

A 0 is returned if the **loadbind** subroutine is successful.

## Error Codes

A -1 is returned if an error is detected, with the **errno** global variable set to an associated error code:

<b>EINVAL</b>	Indicates that either the <i>ImportPointer</i> or <i>ExportPointer</i> parameter is not valid (the pointer to the <i>ExportPointer</i> or <i>ImportPointer</i> parameter does not correspond to a loaded program module or library).
<b>ENOMEM</b>	Indicates that the program requires more memory than allowed by the system-imposed maximum.

After an error is returned by the **loadbind** subroutine, you may also use the **loadquery** subroutine to obtain additional information about the **loadbind** error.

## Related Information

The **load** ("load and loadAndInit Subroutines" on page 727) subroutine, **loadquery** ("loadquery Subroutine" on page 733) subroutine, **unload** subroutine.

The **ld** command.

## loadquery Subroutine

### Purpose

Returns error information from the **load** or **exec** subroutine; also provides a list of object files loaded for the current process.

### Syntax

```
int loadquery( Flags, Buffer, BufferLength)
int Flags;
void *Buffer;
unsigned int BufferLength;
```

### Description

The **loadquery** subroutine obtains detailed information about an error reported on the last **load** or **exec** subroutine executed by a calling process. The **loadquery** subroutine may also be used to obtain a list of object file names for all object files that have been loaded for the current process, or the library path that was used at process exec time.

### Parameters

<i>Buffer</i>	Points to a <i>Buffer</i> in which to store the information.
<i>BufferLength</i>	Specifies the number of bytes available in the <i>Buffer</i> parameter.

*Flags*

Specifies the action of the **loadquery** subroutine as follows:

#### **L\_GETINFO**

Returns a list of all object files loaded for the current process, and stores the list in the *Buffer* parameter. The object file information is contained in a sequence of **LD\_INFO** structures as defined in the **sys/ldr.h** file. Each structure contains the module location in virtual memory and the path name that was used to load it into memory. The file descriptor field in the **LD\_INFO** structure is not filled in by this function.

#### **L\_GETMESSAGE**

Returns detailed error information describing the failure of a previously invoked **load** or **exec** function, and stores the error message information in *Buffer*. Upon successful return from this function the beginning of the *Buffer* contains an array of character pointers. Each character pointer points to a string in the buffer containing a loader error message. The character array ends with a null character pointer. Each error message string consists of an ASCII message number followed by zero or more characters of error-specific message data. Valid message numbers are listed in the **sys/ldr.h** file.

You can format the error messages returned by the **L\_GETMESSAGE** function and write them to standard error using the standard system command **/usr/sbin/execerror** as follows:

```
char *buffer[1024];
buffer[0] = "execerror";
buffer[1] = "name of program that failed to load";
loadquery(L_GETMESSAGES, &buffer[2], \
    sizeof buffer-2*sizeof(char*));
execvp("/usr/sbin/execerror",buffer);
```

This sample code causes the application to terminate after the messages are written to standard error.

#### **L\_GETLIBPATH**

Returns the library path that was used at process exec time. The library path is a null terminated character string.

#### **L\_GETXINFO**

Returns a list of all object files loaded for the current process and stores the list in the *Buffer* parameter. The object file information is contained in a sequence of **LD\_XINFO** structures as defined in the **sys/ldr.h** file. Each structure contains the module location in virtual memory and the path name that was used to load it into memory. The file descriptor field in the **LD\_XINFO** structure is not filled in by this function.

## Return Values

Upon successful completion, **loadquery** returns the requested information in the caller's buffer specified by the *Buffer* and *BufferLength* parameters.

## Error Codes

The **loadquery** subroutine returns with a return code of -1 and the **errno** global variable is set to one of the following when an error condition is detected:

- ENOMEM** Indicates that the caller's buffer specified by the *Buffer* and *BufferLength* parameters is too small to return the information requested. When this occurs, the information in the buffer is undefined.
- EINVAL** Indicates the function specified in the *Flags* parameter is not valid.
- EFAULT** Indicates the address specified in the *Buffer* parameter is not valid.

## Related Information

The **exec** (“exec: execl, execl, execlp, execl, execl, execl, or exect Subroutine” on page 236) subroutine, **load** (“load and loadAndInit Subroutines” on page 727) subroutine, **loadbind** (“loadbind Subroutine” on page 731) subroutine, **unload** subroutine.

The **ld** command.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

The **sys/ldr.h** file.

---

## localeconv Subroutine

### Purpose

Sets the locale-dependent conventions of an object.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <locale.h>
struct lconv *localeconv ( )
```

### Description

The **localeconv** subroutine sets the components of an object using the **lconv** structure. The **lconv** structure contains values appropriate for the formatting of numeric quantities (monetary and otherwise) according to the rules of the current locale.

The fields of the structure with the type **char \*** are strings, any of which (except `decimal_point`) can point to a null string, which indicates that the value is not available in the current locale or is of zero length. The fields with type **char** are nonnegative numbers, any of which can be the **CHAR\_MAX** value which indicates that the value is not available in the current locale. The fields of the **lconv** structure include the following:

`char *decimal_point`  
`char *thousands_sep`

The decimal-point character used to format non-monetary quantities.  
The character used to separate groups of digits to the left of the decimal point in formatted non-monetary quantities.

`char *grouping`

A string whose elements indicate the size of each group of digits in formatted non-monetary quantities.

The value of the `grouping` field is interpreted according to the following:

#### **CHAR\_MAX**

No further grouping is to be performed.

**0** The previous element is to be repeatedly used for the remainder of the digits.

**other** The value is the number of digits that comprise the current group. The next element is examined to determine the size of the next group of digits to the left of the current group.

char *int_curr_symbol	The international currency symbol applicable to the current locale, left-justified within a four-character space-padded field. The character sequences are in accordance with those specified in ISO 4217, "Codes for the Representation of Currency and Funds."
char *currency_symbol	The local currency symbol applicable to the current locale.
char *mon_decimal_point	The decimal point used to format monetary quantities.
char *mon_thousands_sep	The separator for groups of digits to the left of the decimal point in formatted monetary quantities.
char *mon_grouping	A string whose elements indicate the size of each group of digits in formatted monetary quantities.
	The value of the mon_grouping field is interpreted according to the following:
	<b>CHAR_MAX</b>
	No further grouping is to be performed.
	<b>0</b> The previous element is to be repeatedly used for the remainder of the digits.
	<b>other</b> The value is the number of digits that comprise the current group. The next element is examined to determine the size of the next group of digits to the left of the current group.
char *positive_sign	The string used to indicate a nonnegative formatted monetary quantity.
char *negative_sign	The string used to indicate a negative formatted monetary quantity.
char int_frac_digits	The number of fractional digits (those to the right of the decimal point) to be displayed in a formatted monetary quantity.
char p_cs_precedes	Set to 1 if the specified currency symbol (the currency_symbol or int_curr_symbol field) precedes the value for a nonnegative formatted monetary quantity and set to 0 if the specified currency symbol follows the value for a nonnegative formatted monetary quantity.
char p_sep_by_space	Set to 1 if the currency_symbol or int_curr_symbol field is separated by a space from the value for a nonnegative formatted monetary quantity and set to 0 if the currency_symbol or int_curr_symbol field is not separated by a space from the value for a nonnegative formatted monetary quantity.
char n_cs_precedes	Set to 1 if the currency_symbol or int_curr_symbol field precedes the value for a negative formatted monetary quantity and set to 0 if the currency_symbol or int_curr_symbol field follows the value for a negative formatted monetary quantity.
char n_sep_by_space	Set to 1 if the currency_symbol or int_curr_symbol field is separated by a space from the value for a negative formatted monetary quantity and set to 0 if the currency_symbol or int_curr_symbol field is not separated by a space from the value for a negative formatted monetary quantity. Set to 2 if the symbol and the sign string are adjacent and separated by a blank character.
char p_sign_posn	Set to a value indicating the positioning of the positive sign (the positive_sign fields) for nonnegative formatted monetary quantity.



char n\_sign\_posn

Set to a value indicating the positioning of the negative sign (the negative\_sign fields) for a negative formatted monetary quantity.

The values of the p\_sign\_posn and n\_sign\_posn fields are interpreted according to the following definitions:

- 0** Parentheses surround the quantity and the specified currency symbol or international currency symbol.
- 1** The sign string precedes the quantity and the currency symbol or international currency symbol.
- 2** The sign string follows the quantity and currency symbol or international currency symbol.
- 3** The sign string immediately precedes the currency symbol or international currency symbol.
- 4** The sign string immediately follows the currency symbol or international currency symbol.

The following table illustrates the rules that can be used by three countries to format monetary quantities:

Country	Formats
Italy	<b>Positive Format:</b> L.1234 <b>Negative Format:</b> -L.1234 <b>International Format:</b> ITL.1234
Norway	<b>Positive Format:</b> krl.234.56 <b>Negative Format:</b> krl.234.56- <b>International Format:</b> NOK 1.234.56
Switzerland	<b>Positive Format:</b> SFrs.1.234.56 <b>Negative Format:</b> SFrs.1.234.56C <b>International Format:</b> CHF 1.234.56

The following table shows the values of the monetary members of the structure returned by the **localeconv** subroutine for these countries:

struct localeconv	Countries
char *in_curr_symbol	<b>Italy:</b> "ITL." <b>Norway:</b> "NOK" <b>Switzerland:</b> "CHF"

<b>struct localeconv</b>	<b>Countries</b>
char *currency_symbol	<b>Italy:</b> "L." <b>Norway:</b> "kr" <b>Switzerland:</b> "SFrs."
char *mon_decimal_point	<b>Italy:</b> " " <b>Norway:</b> "." <b>Switzerland:</b> " "
char *mon_thousands_sep	<b>Italy:</b> " " <b>Norway:</b> "." <b>Switzerland:</b> " "
char *mon_grouping	<b>Italy:</b> "\3" <b>Norway:</b> "\3" <b>Switzerland:</b> "\3"
char *positive_sign	<b>Italy:</b> " " <b>Norway:</b> " " <b>Switzerland:</b> " "
char *negative_sign	<b>Italy:</b> " _" <b>Norway:</b> " _" <b>Switzerland:</b> "C"
char int_frac_digits	<b>Italy:</b> 0 <b>Norway:</b> 2 <b>Switzerland:</b> 2
char frac_digits	<b>Italy:</b> 0 <b>Norway:</b> 2 <b>Switzerland:</b> 2

<b>struct localeconv</b>	<b>Countries</b>
char p_cs_precedes	<b>Italy:</b> 1 <b>Norway:</b> 1 <b>Switzerland:</b> 1
char p_sep_by_space	<b>Italy:</b> 0 <b>Norway:</b> 0 <b>Switzerland:</b> 0
char n_cs_precedes	<b>Italy:</b> 1 <b>Norway:</b> 1 <b>Switzerland:</b> 1
char n_sep_by_space	<b>Italy:</b> 0 <b>Norway:</b> 0 <b>Switzerland:</b> 0
char p_sign_posn	<b>Italy:</b> 1 <b>Norway:</b> 1 <b>Switzerland:</b> 1
char n_sign_posn	<b>Italy:</b> 1 <b>Norway:</b> 2 <b>Switzerland:</b> 2

## Return Values

A pointer to the filled-in object is returned. In addition, calls to the **setlocale** subroutine with the **LC\_ALL**, **LC\_MONETARY** or **LC\_NUMERIC** categories may cause subsequent calls to the **localeconv** subroutine to return different values based on the selection of the locale.

**Note:** The structure pointed to by the return value is not modified by the program but may be overwritten by a subsequent call to the **localeconv** subroutine.

## Related Information

The “nl\_langinfo Subroutine” on page 907, **rpmatch** subroutine, **setlocale** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

## lockfx, lockf, flock, or lockf64 Subroutine

### Purpose

Locks and unlocks sections of open files.

### Libraries

**lockfx, lockf:** Standard C Library (**libc.a**)

**flock:** Berkeley Compatibility Library (**libbsd.a**)

### Syntax

```
#include <fcntl.h>
```

```
int lockfx (FileDescriptor,  
Command, Argument)  
int FileDescriptor;  
int Command;  
struct flock * Argument;  
#include <sys/lockf.h>  
#include <unistd.h>
```

```
int lockf  
(FileDescriptor, Request, Size)  
int FileDescriptor;  
int Request;  
off_t Size;
```

```
int lockf64 (FileDescriptor,  
Request, Size)  
int FileDescriptor;  
int Request;  
off64_t Size;  
#include <sys/file.h>
```

```
int flock (FileDescriptor, Operation)  
int FileDescriptor;  
int Operation;
```

### Description

**Attention:** Buffered I/O does not work properly when used with file locking. Do not use the standard I/O package routines on files that are going to be locked.

The **lockfx** subroutine locks and unlocks sections of an open file. The **lockfx** subroutine provides a subset of the locking function provided by the **fcntl** (“fcntl, dup, or dup2 Subroutine” on page 255) subroutine.

The **lockf** subroutine also locks and unlocks sections of an open file. However, its interface is limited to setting only write (exclusive) locks.

Although the **lockfx**, **lockf**, **flock**, and **fcntl** interfaces are all different, their implementations are fully integrated. Therefore, locks obtained from one subroutine are honored and enforced by any of the lock subroutines.

The *Operation* parameter to the **lockfx** subroutine, which creates the lock, determines whether it is a read lock or a write lock.

The file descriptor on which a write lock is being placed must have been opened with write access.

**lockf64** is equivalent to **lockf** except that a 64-bit lock request size can be given. For **lockf**, the largest value which can be used is **OFF\_MAX**, for **lockf64**, the largest value is **LONGLONG\_MAX**.

In the large file enabled programming environment, **lockf** is redefined to be **lockf64**.

The **flock** subroutine locks and unlocks entire files. This is a limited interface maintained for BSD compatibility, although its behavior differs from BSD in a few subtle ways. To apply a shared lock, the file must be opened for reading. To apply an exclusive lock, the file must be opened for writing.

Locks are not inherited. Therefore, a child process cannot unlock a file locked by the parent process.

## Parameters

*Argument*

A pointer to a structure of type **flock**, defined in the **flock.h** file.

*Command*

Specifies one of the following constants for the **lockfx** subroutine:

### **F\_SETLK**

Sets or clears a file lock. The *l\_type* field of the **flock** structure indicates whether to establish or remove a read or write lock. If a read or write lock cannot be set, the **lockfx** subroutine returns immediately with an error value of -1.

### **F\_SETLKW**

Performs the same function as **F\_SETLK** unless a read or write lock is blocked by existing locks. In that case, the process sleeps until the section of the file is free to be locked.

### **F\_GETLK**

Gets the first lock that blocks the lock described in the **flock** structure. If a lock is found, the retrieved information overwrites the information in the **flock** structure. If no lock is found that would prevent this lock from being created, the structure is passed back unchanged except that the *l\_type* field is set to **F\_UNLCK**.

*FileDescriptor*

A file descriptor returned by a successful **open** or **fcntl** subroutine, identifying the file to which the lock is to be applied or removed.

*Operation*

Specifies one of the following constants for the **flock** subroutine:

### **LOCK\_SH**

Apply a shared (read) lock.

### **LOCK\_EX**

Apply an exclusive (write) lock.

### **LOCK\_NB**

Do not block when locking. This value can be logically ORed with either **LOCK\_SH** or **LOCK\_EX**.

### **LOCK\_UN**

Remove a lock.

<i>Request</i>	Specifies one of the following constants for the <b>lockf</b> subroutine:
	<b>F_ULOCK</b> Unlocks a previously locked region in the file.
	<b>F_LOCK</b> Locks the region for exclusive (write) use. This request causes the calling process to sleep if the requested region overlaps a locked region, and to resume when granted the lock.
	<b>F_TEST</b> Tests to see if another process has already locked a region. The <b>lockf</b> subroutine returns 0 if the region is unlocked. If the region is locked, then -1 is returned and the <b>errno</b> global variable is set to <b>EACCES</b> .
	<b>F_TLOCK</b> Locks the region for exclusive use if another process has not already locked the region. If the region has already been locked by another process, the <b>lockf</b> subroutine returns a -1 and the <b>errno</b> global variable is set to <b>EACCES</b> .
<i>Size</i>	The number of bytes to be locked or unlocked for the <b>lockf</b> subroutine. The region starts at the current location in the open file, and extends forward if the <i>Size</i> value is positive and backward if the <i>Size</i> value is negative. If the <i>Size</i> value is 0, the region starts at the current location and extends forward to the maximum possible file size, including the unallocated space after the end of the file.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **lockfx**, **lockf**, and **flock** subroutines fail if one of the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid open file descriptor.
<b>EINVAL</b>	The function argument is not one of <b>F_LOCK</b> , <b>F_TLOCK</b> , <b>F_TEST</b> or <b>F_ULOCK</b> ; or <i>size</i> plus the current file offset is less than 0.
<b>EINVAL</b>	An attempt was made to lock a fifo or pipe.
<b>EDEADLK</b>	The lock is blocked by a lock from another process. Putting the calling process to sleep while waiting for the other lock to become free would cause a deadlock.
<b>ENOLCK</b>	The lock table is full. Too many regions are already locked.
<b>EINTR</b>	The command parameter was <b>F_SETLKW</b> and the process received a signal while waiting to acquire the lock.
<b>EOVERFLOW</b>	The offset of the first, or if <i>size</i> is not 0 then the last, byte in the requested section cannot be represented correctly in an object of type <i>off_t</i> .

The **lockfx** and **lockf** subroutines fail if one of the following is true:

<b>EACCES</b>	The <i>Command</i> parameter is <b>F_SETLK</b> , the <i>l_type</i> field is <b>F_RDLCK</b> , and the segment of the file to be locked is already write-locked by another process.
<b>EACCES</b>	The <i>Command</i> parameter is <b>F_SETLK</b> , the <i>l_type</i> field is <b>F_WRLCK</b> , and the segment of a file to be locked is already read-locked or write-locked by another process.

The **flock** subroutine fails if the following is true:

<b>EWOULDBLOCK</b>	The file is locked and the <b>LOCK_NB</b> option was specified.
--------------------	---

## Related Information

The **close** (“close Subroutine” on page 175) subroutine, **exec**: **execl**, **execv**, **execle**, **execlp**, **execvp**, or **exec** (“exec: execl, execle, execlp, execv, execve, execvp, or exec Subroutine” on page 236) subroutine, **fcntl** (“fcntl, dup, or dup2 Subroutine” on page 255) subroutine, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **open**, **openx**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## log10, log10f, or log10l Subroutine

### Purpose

Computes the Base 10 logarithm.

### Syntax

```
#include <math.h>
```

```
float log10f (x)
float x;
```

```
long double log10l (x)
long double x;
```

```
double log10 (x)
double x;
```

### Description

The **log10f**, **log10l**, and **log10** subroutines compute the base 10 logarithm of the *x* parameter,  $\log_{10}(x)$ .

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Specifies the value to be computed.

### Return Values

Upon successful completion, the **log10**, **log10f**, and **log10l** subroutines return the base 10 logarithm of *x*.

If *x* is  $\pm 0$ , a pole error occurs and **log10**, **log10f**, and **log10l** subroutines return **-HUGE\_VAL**, **-HUGE\_VALF** and **-HUGE\_VALL**, respectively.

For finite values of *x* that are less than 0, or if *x* is  $-\text{Inf}$ , a domain error occurs, and a NaN is returned.

If *x* is NaN, a NaN is returned.

If *x* is 1, +0 is returned.

If *x* is  $+\text{Inf}$ ,  $+\text{Inf}$  is returned.

## Error Codes

When using the **libm.a** library:

**log10** If the *x* parameter is less than 0, the **log10** subroutine returns a **NaNQ** value and sets **errno** to **EDOM**. If *x* = 0, the **log10** subroutine returns a **-HUGE\_VAL** value and sets **errno** to **ERANGE**.

When using **libmsaa.a(-lmsaa)**:

**log10** If the *x* parameter is not positive, the **log10** subroutine returns a **-HUGE\_VAL** value and sets **errno** to **EDOM**. A message indicating DOMAIN error (or SING error when *x* = 0) is output to standard error.

**log10** If *x* < 0, **log10l** returns the value **NaNQ** and sets **errno** to **EDOM**. If *x* equals 0, **log10l** returns the value **-HUGE\_VAL** but does not modify **errno**.

## Related Information

“feclearexcept Subroutine” on page 265, “fetetestexcept Subroutine” on page 273, “class, \_class, finite, isnan, or unordered Subroutines” on page 167, and “madd, msub, mult, mdiv, pow, gcd, invert, rpow, msqrt, mcmp, move, min, omin, fmin, m\_in, mout, omout, fmout, m\_out, sdiv, or itom Subroutine” on page 785.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## log1pf, log1pl, or log1p Subroutine

### Purpose

Computes a natural logarithm.

### Syntax

```
#include <math.h>
```

```
float log1pf (x)  
float x;
```

```
long double log1pl (x)  
long double x;
```

```
double log1p (x)  
double x;
```

### Description

The **log1pf**, **log1pl**, and **log1p** subroutines compute  $\log_e (1.0 + x)$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Specifies the value to be computed.



## Return Values

Upon successful completion, the **log1pf**, **log1pl**, and **log1p** subroutines return the natural logarithm of  $1.0 + x$ .

If  $x$  is  $-1$ , a pole error occurs and the **log1pf**, **log1pl**, and **log1p** subroutines return **-HUGE\_VALF**, **-HUGE\_VALL**, and **-HUGE\_VAL**, respectively.

For finite values of  $x$  that are less than  $-1$ , or if  $x$  is  $-\text{Inf}$ , a domain error occurs, and a NaN is returned.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$ , or  $+\text{Inf}$ ,  $x$  is returned.

If  $x$  is subnormal, a range error may occur and  $x$  should be returned.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## log2, log2f, or log2l Subroutine

### Purpose

Computes base 2 logarithm.

### Syntax

```
#include <math.h>
```

```
double log2 (x)
double x;
```

```
float log2f (x)
float x;
```

```
long double log2l (x)
long double x;
```

### Description

The **log2**, **log2f**, and **log2l** subroutines compute the base 2 logarithm of the  $x$  parameter,  $\log_2(x)$ .

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

## Return Values

Upon successful completion, the **log2**, **log2f**, and **log2l** subroutines return the base 2 logarithm of  $x$ .

If  $x$  is  $\pm 0$ , a pole error occurs and the **log2**, **log2f**, and **log2l** subroutines return **-HUGE\_VAL**, **-HUGE\_VALF**, and **-HUGE\_VALL**, respectively.

For finite values of  $x$  that are less than 0, or if  $x$  is **-Inf**, a domain error occurs, and a NaN is returned.

If  $x$  is NaN, a NaN is returned.

If  $x$  is 1, **+0** is returned.

If  $x$  is **+Inf**,  $x$  is returned.

## Related Information

“**feclearexcept** Subroutine” on page 265 and “**fetestexcept** Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## logbd32, logbd64, and logbd128 Subroutines

### Purpose

Computes the radix-independent exponent.

### Syntax

```
#include <math.h>
```

```
_Decimal132 logbd32 (x)  
_Decimal132 x;
```

```
_Decimal164 logbd64 (x)  
_Decimal164 x;
```

```
_Decimal128 logbd128 (x)  
_Decimal128 x;
```

### Description

The **logbd32**, **logbd64**, and **logbd128** subroutines compute the exponent of  $x$ , which is the integral part of  $\log_r |x|$ , as a signed floating-point value, for nonzero  $x$ , where  $r$  is the radix of the machine’s decimal floating-point arithmetic. For AIX, FLT\_RADIX  $r=10$ .

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

### Return Values

Upon successful completion, the **logbd32**, **logbd64**, and **logbd128** subroutines return the exponent of  $x$ .

If  $x$  is  $\pm 0$ , a pole error occurs and the **logbd32**, **logbd64**, and **logbd128** subroutines return **-HUGE\_VAL\_D32**, **-HUGE\_VAL\_D64**, and **-HUGE\_VAL\_D128**, respectively.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm\text{Inf}$ ,  $+\text{Inf}$  is returned.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## ilogbd32, ilogbd64, and ilogbd128 Subroutines

### Purpose

Returns an unbiased exponent.

### Syntax

```
#include <math.h>
```

```
int ilogbd32 (x)
   _Decimal32 x;
```

```
int ilogbd64 (x)
   _Decimal64 x;
```

```
int ilogbd128 (x)
   _Decimal128 x;
```

### Description

The **ilogbd32**, **ilogbd64**, and **ilogbd128** subroutines return the integral part of  $\log_r |x|$  as a signed integral value, for nonzero  $x$ , where  $r$  is the radix of the machine's floating-point arithmetic ( $r=10$ ).

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$                     Specifies the value to be computed.

### Return Values

Upon successful completion, the **ilogbd32**, **ilogbd64**, and **ilogbd128** subroutines return the exponent part of  $x$  as a signed integer value. They are equivalent to calling the corresponding **logb** functions and casting the returned value to type **int**.

If  $x$  is 0, a domain error occurs, and the value **FP\_ILOGB0** is returned.

If  $x$  is  $\pm\text{Inf}$ , a domain error occurs, and the value **{INT\_MAX}** is returned.

If  $x$  is a NaN, a domain error occurs, and the value **FP\_ILOGBNAN** is returned.

If the correct value is greater than **{INT\_MAX}**, **{INT\_MAX}** is returned and a domain error occurs.

If the correct value is less than **{INT\_MIN}**, **{INT\_MIN}** is returned and a domain error occurs.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

`math.h` in *AIX 5L Version 5.3 Files Reference*.

---

## logbf, logbl, or logb Subroutine

### Purpose

Computes the radix-independent exponent.

### Syntax

```
#include <math.h>
```

```
float logbf (x)
float x;
```

```
long double logbl (x)
long double x;
```

```
double logb(x)
double x;
```

### Description

The **logbf** and **logbl** subroutines compute the exponent of  $x$ , which is the integral part of  $\log_r |x|$ , as a signed floating-point value, for nonzero  $x$ , where  $r$  is the radix of the machine's floating-point arithmetic. For AIX, FLT\_RADIX  $r=2$ .

If  $x$  is subnormal, it is treated as though it were normalized; thus for finite positive  $x$ :

$$1 \leq x * FLT\_RADIX^{-\text{logb}(x)} < FLT\_RADIX$$

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

**Note:** When the  $x$  parameter is finite and not zero, the **logb** ( $x$ ) subroutine satisfies the following equation:

$$1 \leq \text{scalb}(|x|, -(\text{int}) \text{logb}(x)) < 2$$

### Parameters

$x$  Specifies the value to be computed.

### Return Values

Upon successful completion, the **logbf** and **logbl** subroutines return the exponent of  $x$ .

If  $x$  is  $\pm 0$ , a pole error occurs and the **logbf** and **logbl** subroutines return **-HUGE\_VALF** and **-HUGE\_VALL**, respectively.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm\text{Inf}$ ,  $+\text{Inf}$  is returned.

## Error Codes

The **logb** function returns **-HUGE\_VAL** when the *x* parameter is set to a value of 0 and sets **errno** to **EDOM**.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## log, logf, or logl Subroutine

### Purpose

Computes the natural logarithm.

### Syntax

```
#include <math.h>
```

```
float logf (x)  
float x;
```

```
long double logl (x)  
long double x;
```

```
double log (x)  
double x;
```

### Description

The **logf**, **logl**, and **log** subroutines compute the natural logarithm of the *x* parameter,  $\log_e(x)$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Specifies the value to be computed.

### Return Values

Upon successful completion, the **logf**, **logl**, and **log** subroutines return the natural logarithm of *x*.

If *x* is  $\pm 0$ , a pole error occurs and the **logf**, **logl**, and **log** subroutines return **-HUGE\_VALF** and **-HUGE\_VALL**, and **-HUGE\_VAL**, respectively.

For finite values of *x* that are less than 0, or if *x* is  $-\text{Inf}$ , a domain error occurs, and a NaN is returned.

If *x* is NaN, a NaN is returned.

If *x* is 1, +0 is returned.

If *x* is  $+\text{Inf}$ , *x* is returned.

## Error Codes

When using the **libm.a** library:

**log** If the *x* parameter is less than 0, the **log** subroutine returns a **NaNQ** value and sets **errno** to **EDOM**. If *x* = 0, the **log** subroutine returns a **-HUGE\_VAL** value but does not modify **errno**.

When using **libmsaa.a(-lmsaa)**:

**log** If the *x* parameter is not positive, the **log** subroutine returns a **-HUGE\_VAL** value, and sets **errno** to a **EDOM** value. A message indicating DOMAIN error (or SING error when *x* = 0) is output to standard error.

**logl** If *x* < 0, the **logl** subroutine returns a **NaNQ** value

## Related Information

“exp, expf, or expl Subroutine” on page 245, “feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, “class, \_class, finite, isnan, or unordered Subroutines” on page 167, and “log10, log10f, or log10l Subroutine” on page 743.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## loginfailed Subroutine

### Purpose

Records an unsuccessful login attempt.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
int loginfailed ( User, Host, Tty, Reason)
char *User;
char *Host;
char *Tty;
int Reason;
```

**Note:** This subroutine is not thread-safe.

### Description

The **loginfailed** subroutine performs the processing necessary when an unsuccessful login attempt occurs. If the specified user name is not valid, the **UNKNOWN\_USER** value is substituted for the user name. This substitution prevents passwords entered as the user name from appearing on screen.

The following attributes in **/etc/security/lastlog** file are updated for the specified user, if the user name is valid:

<b>time_last_unsuccessful_login</b>	Contains the current time.
<b>tty_last_unsuccessful_login</b>	Contains the value specified by the <i>Tty</i> parameter.
<b>host_last_unsuccessful_login</b>	Contains the value specified by the <i>Host</i> parameter, or the local hostname if the <i>Host</i> parameter is a null value.

## **unsuccessful\_login\_count**

Indicates the number of unsuccessful login attempts. The **loginfailed** subroutine increments this attribute by one for each failed attempt.

A login failure audit record is cut to indicate that an unsuccessful login attempt occurred. A **utmp** entry is appended to **/etc/security/failedlogin** file, which tracks all failed login attempts.

If the current unsuccessful login and the previously recorded unsuccessful logins constitute too many unsuccessful login attempts within too short of a time period (as specified by the **logindisable** and **logininterval** port attributes), the port is locked. When a port is locked, a **PORT\_Locked** audit record is written to inform the system administrator that the port has been locked.

If the login retry delay is enabled (as specified by the **logindelay** port attribute), a sleep occurs before this subroutine returns. The length of the sleep (in seconds) is determined by the **logindelay** value multiplied by the number of unsuccessful login attempts that occurred in this process.

## **Parameters**

<i>User</i>	Specifies the user's login name who has unsuccessfully attempted to login.
<i>Host</i>	Specifies the name of the host from which the user attempted to login. If the <i>Host</i> parameter is Null, the name of the local host is used.
<i>Tty</i>	Specifies the name of the terminal on which the user attempted to login.
<i>Reason</i>	Specifies a reason code for the login failure. Valid values are <b>AUDIT_FAIL</b> and <b>AUDIT_FAIL_AUTH</b> defined in the <b>sys/audit.h</b> file.

## **Security**

Access Control: The calling process must have access to the account information in the user database and the port information in the port database.

File Accessed:

<b>Mode</b>	<b>File</b>
<b>r</b>	<b>/etc/security/user</b>
<b>rw</b>	<b>/etc/security/lastlog</b>
<b>r</b>	<b>/etc/security/login.cfg</b>
<b>rw</b>	<b>/etc/security/portlog</b>
<b>w</b>	<b>/etc/security/failedlogin</b>

Auditing Events:

<b>Event</b>	<b>Information</b>
<b>USER_Login</b>	username
<b>PORT_Locked</b>	portname

## **Return Values**

Upon successful completion, the **loginfailed** subroutine returns a value of 0. If an error occurs, a value of -1 is returned and **errno** is set to indicate the error.

## **Error Codes**

The **loginfailed** subroutine fails if one or more of the following values is true:

**EACCESS** The current process does not have access to the user or port database.

**EPERM** The current process does not have permission to write an audit record.

## Related Information

The **authenticate** (“authenticate Subroutine” on page 113) subroutine, **getpcred** (“getpcred Subroutine” on page 403) subroutine, **getpenv** (“getpenv Subroutine” on page 405) subroutine, **loginrestrictions** (“loginrestrictions Subroutine”) subroutine, **loginsuccess** (“loginsuccess Subroutine” on page 757) subroutine, **setpcred** subroutine, **setpenv** subroutine.

List of Security and Auditing Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## loginrestrictions Subroutine

### Purpose

Determines if a user is allowed to access the system.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
#include <login.h>
```

```
int loginrestrictions (Name, Mode, Tty, Msg)
char * Name;
int Mode;
char * Tty;
char ** Msg;
```

**Note:** This subroutine is not thread-safe.

### Description

The **loginrestrictions** subroutine determines if the user specified by the *Name* parameter is allowed to access the system. The *Mode* parameter gives the mode of account usage and the *Tty* parameter defines the terminal used for access. The *Msg* parameter returns an informational message explaining why the **loginrestrictions** subroutine failed.

This subroutine is unsuccessful if any of the following conditions exists:

- The user’s account has expired as defined by the **expires** user attribute.
- The user’s account has been locked as defined by the **account\_locked** user attribute.
- The user attempted too many unsuccessful logins as defined by the **loginretries** user attribute.
- The user is not allowed to access the given terminal as defined by the **ttys** user attribute.
- The user is not allowed to access the system at the present time as defined by the **logintimes** user attribute.
- The *Mode* parameter is set to the **S\_LOGIN** value or the **S\_RLOGIN** value, and too many users are logged in as defined by the **maxlogins** system attribute.



- The *Mode* parameter is set to the **S\_LOGIN** value and the user is not allowed to log in as defined by the **login** user attribute.
- The *Mode* parameter is set to the **S\_RLOGIN** value and the user is not allowed to log in from the network as defined by the **rlogin** user attribute.
- The *Mode* parameter is set to the **S\_SU** value and other users are not allowed to use the **su** command as defined by the **su** user attribute, or the group ID of the current process cannot use the **su** command to switch to this user as defined by the **sugroups** user attribute.
- The *Mode* parameter is set to the **S\_DAEMON** value and the user is not allowed to run processes from the **cron** or **src** subsystem as defined by the **daemon** user attribute.
- The terminal is locked as defined by the **locktime** port attribute.
- The user cannot use the terminal to access the system at the present time as defined by the **logintimes** port attribute.
- The user is not the root user and the **/etc/nologin** file exists.

**Note:** The **loginrestrictions** subroutine is not safe in a multi-threaded environment. To use **loginrestrictions** in a threaded application, the application must keep the integrity of each thread.

## Parameters

<i>Name</i>	Specifies the user's login name whose account is to be validated.
<i>Mode</i>	Specifies the mode of usage. Valid values as defined in the <b>login.h</b> file are listed below. The <i>Mode</i> parameter has a value of 0 or one of the following values:
	<b>S_LOGIN</b> Verifies that local logins are permitted for this account.
	<b>S_SU</b> Verifies that the <b>su</b> command is permitted and the current process has a group ID that can invoke the <b>su</b> command to switch to the account.
	<b>S_DAEMON</b> Verifies the account can invoke daemon or batch programs through the <b>src</b> or <b>cron</b> subsystems.
	<b>S_RLOGIN</b> Verifies the account can be used for remote logins through the <b>rlogind</b> or <b>telnetd</b> programs.
<i>Tty</i>	Specifies the terminal of the originating activity. If this parameter is a null pointer or a null string, no tty origin checking is done.
<i>Msg</i>	Returns an informative message indicating why the <b>loginrestrictions</b> subroutine failed. Upon return, the value is either a pointer to a valid string within memory allocated storage or a null value. If a message is displayed, it is provided based on the user interface.

## Security

**Access Control:**The calling process must have access to the account information in the user database and the port information in the port database.

File Accessed:

<b>Mode</b>	<b>Files</b>
r	/etc/security/user
r	/etc/security/login.cfg
r	/etc/security/portlog
r	/etc/passwd

## Return Values

If the account is valid for the specified usage, the **loginrestrictions** subroutine returns a value of 0. Otherwise, a value of -1 is returned, the **errno** global value is set to the appropriate error code, and the *Msg* parameter returns an informative message explaining why the specified account usage is invalid.

## Error Codes

The **loginrestrictions** subroutine fails if one or more of the following values is true:

<b>ENOENT</b>	The user specified does not have an account.
<b>ESTALE</b>	The user's account is expired.
<b>EPERM</b>	The user's account is locked, the specified terminal is locked, the user has had too many unsuccessful login attempts, or the user cannot log in because the <b>/etc/nologin</b> file exists.
<b>EACCES</b>	One of the following conditions exists: <ul style="list-style-type: none"><li>• The specified terminal does not have access to the specified account.</li><li>• The <i>Mode</i> parameter is the <b>S_SU</b> value and the current process is not permitted to use the <b>su</b> command to access the specified user.</li><li>• Access to the account is not permitted in the specified mode.</li><li>• Access to the account is not permitted at the current time.</li><li>• Access to the system with the specified terminal is not permitted at the current time.</li></ul>
<b>EAGAIN</b>	The <i>Mode</i> parameter is either the <b>S_LOGIN</b> value or the <b>S_RLOGIN</b> value, and all the user licenses are in use.
<b>EINVAL</b>	The <i>Mode</i> parameter has a value other than <b>S_LOGIN</b> , <b>S_SU</b> , <b>S_DAEMON</b> , <b>S_RLOGIN</b> , or <b>0</b> .

## Related Information

The **authenticate** (“authenticate Subroutine” on page 113) subroutine, **getpcred** (“getpcred Subroutine” on page 403) subroutine, **getpenv** (“getpenv Subroutine” on page 405) subroutine, **loginfailed** (“loginfailed Subroutine” on page 750) subroutine, **loginsuccess** (“loginsuccess Subroutine” on page 757) subroutine, **setpcred** subroutine, **setpenv** subroutine.

The **cron** daemon.

The **login** command, **rlogin** command, **telnet**, **tn**, or **tn3270** command, **su** command.

List of Security and Auditing Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## loginrestrictionsx Subroutine

### Purpose

Determines, in multiple methods, if a user is allowed to access the system.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
#include <login.h>
```

```

int loginrestrictionsx (Name, Mode, Tty, Message, State)
char * Name;
int Mode;
char *Tty;
char **Message;
char **State;

```

## Description

The **loginrestrictionsx** subroutine determines if the user specified by the *Name* parameter is allowed to access the system. The *Mode* parameter gives the mode of account usage, and the *Tty* parameter defines the terminal used for access. The *Msg* parameter returns an informational message explaining why the **loginrestrictionsx** subroutine failed. The user's **SYSTEM** attribute determines the administrative domains to examine for permission.

The *State* parameter contains information about the login restrictions for the user. A call to the **authenticatex** subroutine will not use an administrative domain for authentication if an earlier call to **loginrestrictionsx** indicated that the user was unable to log in using that administrative domain's authentication data. The result is that administrative domains that are used for authentication must permit the user to log in. The *State* parameter returned by **loginrestrictionsx** can be used as input to a subsequent call to the **authenticatex** subroutine.

This subroutine is unsuccessful if any of the following conditions exists:

- The user's account has been locked as defined by the **account\_locked** user attribute.
- The user's account has expired as defined by the **expires** user attribute.
- The *Mode* parameter is set to the **S\_LOGIN** value or the **S\_RLOGIN** value, and too many users are logged in as defined by the **maxlogins** system attribute.
- The *Mode* parameter is not set to the **S\_SU** or **S\_DAEMON** value, and the user is not allowed to log in to the current host as defined by the user's **hostallowedlogin** and **hostdeniedlogin** attributes.
- The user is not allowed to access the system at the present time as defined by the **logintimes** user attribute.
- The user attempted too many unsuccessful logins as defined by the **loginretries** user attribute.
- The user is not allowed to access the given terminal or network protocol as defined by the **ttys** user attribute. This test is not performed when the *Mode* parameter is set to the **S\_DAEMON** value.
- The *Mode* parameter is set to the **S\_LOGIN** value, and the user is not allowed to log in as defined by the **login** user attribute.
- The *Mode* parameter is set to the **S\_RLOGIN** value and the user is not allowed to log in from the network as defined by the **rlogin** user attribute.
- The *Mode* parameter is set to the **S\_SU** value, and other users are not allowed to use the **su** command as defined by the **su** user attribute; or, the group ID of the current process cannot use the **su** command to switch to this user as defined by the **sugroups** user attribute.
- The *Mode* parameter is set to the **S\_DAEMON** value, and the user is not allowed to run processes from the **cron** or **src** subsystem as defined by the **daemon** user attribute.
- The terminal is locked as defined by the **locktime** port attribute.
- The user cannot use the terminal to access the system at the present time as defined by the **logintimes** port attribute.
- The user is not the root user, and the **/etc/nologin** file exists.

Additional restrictions can be enforced by loadable authentication modules for any administrative domain used in the user's **SYSTEM** attribute.

## Parameters

<i>Name</i>	Specifies the user's login name whose account is to be validated.
<i>Mode</i>	Specifies the mode of usage. The valid values in the following list are defined in the <b>login.h</b> file. The <i>Mode</i> parameter has a value of 0 or one of the following values:  <b>S_LOGIN</b> Verifies that local logins are permitted for this account.  <b>S_SU</b> Verifies that the <b>su</b> command is permitted and the current process has a group ID that can invoke the <b>su</b> command to switch to the account.  <b>S_DAEMON</b> Verifies that the account can invoke daemon or batch programs through the <b>src</b> or <b>cron</b> subsystems.  <b>S_RLOGIN</b> Verifies that the account can be used for remote logins through the <b>rlogind</b> or <b>telnetd</b> programs.
<i>Tty</i>	Specifies the terminal of the originating activity. If this parameter is a null pointer or a null string, no <b>tty</b> origin checking is done. The <i>Tty</i> parameter can also have the value <b>RSH</b> or <b>REXEC</b> to indicate that the caller is the <b>rsh</b> or <b>rexec</b> command.
<i>Message</i>	Returns an informative message indicating why the <b>loginrestrictionsx</b> subroutine failed. Upon return, the value is either a pointer to a valid string within memory-allocated storage or a null value. If a message is displayed, it is provided based on the user interface.
<i>State</i>	Points to a pointer that the <b>loginrestrictionsx</b> subroutine allocates memory for and fills in. The <i>State</i> parameter can also be the result of an earlier call to the <b>authenticatex</b> subroutine. The <i>State</i> parameter contains information about the results of the <b>loginrestrictionsx</b> subroutine for each term in the user's <b>SYSTEM</b> attribute. The calling application is responsible for freeing this memory when it is no longer needed for a subsequent call to the <b>authenticatex</b> , <b>passwdexpiredx</b> , or <b>chpasswdx</b> subroutines.

## Security

Access Control: The calling process must have access to the account information in the user database and the port information in the port database.

Files accessed:

Mode	File
r	/etc/security/user
r	/etc/security/login.cfg
r	/etc/security/portlog
r	/etc/passwd

## Return Values

If the account is valid for the specified usage, the **loginrestrictionsx** subroutine returns a value of 0. Otherwise, a value of -1 is returned, the **errno** global value is set to the appropriate error code, and the *Message* parameter returns an informative message explaining why the specified account usage is invalid.

## Error Codes

If the `loginrestrictionsx` subroutine fails if one of the following values is true:

<b>EACCES</b>	One of the following conditions exists: <ul style="list-style-type: none"><li>• The specified terminal does not have access to the specified account.</li><li>• The <i>Mode</i> parameter is the <b>S_SU</b> value, and the current process is not permitted to use the <b>su</b> command to access the specified user.</li><li>• Access to the account is not permitted in the specified mode.</li><li>• Access to the account is not permitted at the current time.</li><li>• Access to the system with the specified terminal is not permitted at the current time.</li></ul>
<b>EAGAIN</b>	The <i>Mode</i> parameter is either the <b>S_LOGIN</b> value or the <b>S_RLOGIN</b> value, and all the user licenses are in use.
<b>EINVAL</b>	The <i>Mode</i> parameter has a value other than <b>S_LOGIN</b> , <b>S_SU</b> , <b>S_DAEMON</b> , <b>S_RLOGIN</b> , or 0.
<b>ENOENT</b>	The user specified does not have an account.
<b>EPERM</b>	The user's account is locked, the specified terminal is locked, the user has had too many unsuccessful login attempts, or the user cannot log in because the <code>/etc/nologin</code> file exists.
<b>ESTALE</b>	The user's account is expired.

## Related Information

The “`authenticatex` Subroutine” on page 115, “`getpcred` Subroutine” on page 403, “`getpenv` Subroutine” on page 405, “`loginfailed` Subroutine” on page 750, “`loginsuccess` Subroutine,” “`getgroupattrs` Subroutine” on page 378, “`getuserpw`, `putuserpw`, or `putuserpwhist` Subroutine” on page 468, `setpcred` Subroutines, `setpenv` Subroutine.

The cron Daemon.

The login Command, rlogin Command, telnet, tn, or tn3270 Command, suCommand.

List of Security and Auditing Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## loginsuccess Subroutine

### Purpose

Records a successful log in.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
int loginsuccess (User, Host, Tty, Msg)
char * User;
char * Host;
char * Tty;
char ** Msg;
```

**Note:** This subroutine is not thread-safe.

## Description

The **loginsuccess** subroutine performs the processing necessary when a user successfully logs into the system. This subroutine updates the following attributes in the **/etc/security/lastlog** file for the specified user:

<b>time_last_login</b>	Contains the current time.
<b>tty_last_login</b>	Contains the value specified by the <i>Tty</i> parameter.
<b>host_last_login</b>	Contains the value specified by the <i>Host</i> parameter or the local host name if the <i>Host</i> parameter is a null value.
<b>unsuccessful_login_count</b>	Indicates the number of unsuccessful login attempts. The <b>loginsuccess</b> subroutine resets this attribute to a value of 0.

Additionally, a login success audit record is cut to indicate in the audit trail that this user has successfully logged in.

A message is returned in the *Msg* parameter that indicates the time, host, and port of the last successful and unsuccessful login. The number of unsuccessful login attempts since the last successful login is also provided to the user.

## Parameters

<i>User</i>	Specifies the login name of the user who has successfully logged in.
<i>Host</i>	Specifies the name of the host from which the user logged in. If the <i>Host</i> parameter is a null value, the name of the local host is used.
<i>Tty</i>	Specifies the name of the terminal which the user used to log in.
<i>Msg</i>	Returns a message indicating the delete time, host, and port of the last successful and unsuccessful logins. The number of unsuccessful login attempts since the last successful login is also provided. Upon return, the value is either a pointer to a valid string within memory allocated storage or a null pointer. It is the responsibility of the calling program to <b>free( )</b> the returned storage.

## Security

**Access Control:** The calling process must have access to the account information in the user database.

File Accessed:

Mode	File
rw	/etc/security/lastlog

Auditing Events:

Event	Information
USER_Login	username

## Return Values

Upon successful completion, the **loginsuccess** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global value is set to indicate the error.

## Error Codes

The **loginsuccess** subroutine fails if one or more of the following values is true:

<b>ENOENT</b>	The specified user does not exist.
<b>EACCES</b>	The current process does not have write access to the user database.
<b>EPERM</b>	The current process does not have permission to write an audit record.

## Related Information

The **authenticate** (“authenticate Subroutine” on page 113) subroutine, **getpcred** (“getpcred Subroutine” on page 403) subroutine, **getpenv** (“getpenv Subroutine” on page 405) subroutine, **loginfailed** (“loginfailed Subroutine” on page 750) subroutine, **loginrestrictions** (“loginrestrictions Subroutine” on page 752) subroutine, **setpcred** subroutine, **setpenv** subroutine.

List of Security and Auditing Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## lpar\_get\_info Subroutine

### Purpose

Retrieves the calling partition’s characteristics.

### Syntax

```
#include <sys/dr.h>

int lpar_get_info (command, lparinfo, bufsize)
int command;
void *lparinfo;
size_t bufsize;
```

### Description

The **lpar\_get\_info** subroutine retrieves LPAR and Micro-Partitioning attributes of both low-frequency use and high-frequency use. Because the low-frequency attributes, as defined in the **lpar\_info\_format1** structure, are static in nature, a reboot is required to effect any change. The high-frequency attributes, as defined in the **lpar\_info\_format2** structure, can be changed dynamically at any time either by the platform or through DLPAR procedures. The latter provides a mechanism for notifying applications of changes. The signature of this system call, its parameter types, and the order of the member fields in both the **lpar\_info\_format1** and **lpar\_info\_format2** structures are specific to the AIX platform.

To see the complete structures of **lpar\_info\_format1** and **lpar\_info\_format2**, see the **dr.h** header file.

The **lpar\_get\_info** system call provides information about the operating system environment, including the following:

- Type of partition: dedicated processor partition or micro-partition
- Type of micro-partition: capped or uncapped
- Variable capacity weight of micro-partition
- Partition name and number
- SMT-capable partition
- SMT-enabled partition

- Minimum, desired, online, and maximum number of virtual processors
- Minimum, online, and maximum number of logical processors
- Minimum, desired, online, and maximum entitled processor capacity
- Minimum, desired, online (megabytes), and maximum number of logical memory blocks (LMBs)
- Maximum number of potential installed physical processors in the server, including unlicensed and potentially hot-pluggable
- Number of active licensed installed physical processors in the server
- Number of processors in the shared processor pool

This subroutine is used by the DRM to determine whether a client partition is migration capable and MSP capable. The kernel presents these capabilities based on the presence of the **hcall-vasi** function set and the type of partition that is evident. If the partition is a VIOS partition, the MSP capability will be noted. Otherwise, the OS partition migration capability will be noted.

## Parameters

<i>command</i>	Specifies whether the user wants <b>format1</b> or <b>format2</b> details.
<i>lparinfo</i>	Pointer to the user-allocated buffer that is passed in.
<i>bufsize</i>	Size of the structure that is passed in.

## Return Values

Upon success, the **lpar\_get\_info** subroutine returns a value of 0. Upon failure, a value of -1 is returned, and **errno** is set to indicate the appropriate error.

## Error Codes

<b>EFAULT</b>	Buffer size is smaller than expected.
<b>EINVAL</b>	Invalid input parameter.
<b>ENOTSUP</b>	The platform does not support this operation.

## Related Information

The `klpar_get_info` kernel service.

---

## lpar\_set\_resources Subroutine

### Purpose

Modifies the calling partition's characteristics.

### Library

Standard C Library (**lib.c**)

### Syntax

```
#include <sys/dr.h>
int lpar_set_resources ( lpar_resource_id, lpar_resource )
int lpar_resource_id;
void *lpar_resource;
```



## Description

The **lpar\_set\_resources** subroutine modifies the configuration attributes (dynamic resources) on a current partition indicated by the *lpar\_resource\_id*. The pointer to a value of the dynamic resource indicated by *lpar\_resource\_id* is passed to this call in *lpar\_resource*. This subroutine modifies one partition dynamic resource at a time. To reconfigure multiple resources, multiple calls must be made. The following resources for the calling partition can be modified:

- Entitled Capacity
- Variable Capacity Weight
- Number of online virtual CPUs
- Number of available memory in MB

These resource IDs are defined in the **<sys/dr.h>** header file. In order to modify the Entitled Capacity and Variable Capacity Weight attributes, ensure that the current partition is an SPLPAR partition; otherwise, an error is returned. The **lpar\_set\_resources** subroutine can only be called in a process owned by a root user (super user) or a user with the **CAP\_EWLM\_AGENT** capability; otherwise, an error is returned.

## Parameters

<i>lpar_resource_id</i>	Identifies the dynamic resource whose value is being changed.
<i>lpar_resource</i>	Pointer to a new value of the dynamic resource identified by the <i>lpar_resource_id</i> .

## Security

The **lpar\_set\_resources** subroutine can only be called in a process owned by a root user (super user) or a user with the **CAP\_EWLM\_AGENT** capability.

## Return Values

Upon success, the **lpar\_set\_resources** subroutine returns a value of 0. Upon failure, a negative value is returned, and **errno** is set to the appropriate error.

## Error Codes

<b>EINVAL</b>	Invalid configuration parameters.
<b>EPERM</b>	Insufficient authority.
<b>EEXIST</b>	Resource already exists.
<b>EBUSY</b>	Resource is busy.
<b>EAGAIN</b>	Resource is temporarily unavailable.
<b>ENOMEM</b>	Resource allocation failed.
<b>ENOTREADY</b>	Resource is not ready.
<b>ENOTSUP</b>	Operation is not supported.
<b>EFAULT/EIO</b>	Operation failed because of an I/O error.
<b>EINPROGRESS</b>	Operation in progress.
<b>ENXIO</b>	Resource is not available.
<b>ERANGE</b>	Parameter value is out of range.
All others	Internal error.

---

## Irint, Irintf, Irintl, Irintd32, Irintd64, and Irintd128 Subroutines

### Purpose

Round to nearest integer value using the current rounding direction.

## Syntax

```
#include <math.h>
```

```
long lrint (x)  
double x;
```

```
long lrintf (x)  
float x;
```

```
long lrintl (x)  
long double x;
```

```
long lrintd32 (x)  
_Decimal32 x;
```

```
long lrintd64 (x)  
_Decimal64 x;
```

```
long lrintd128 (x)  
_Decimal128 x;
```

## Description

The **lrint**, **lrintf**, **lrintl**, **lrintd32**, **lrintd64**, and **lrintd128** subroutines round the *x* parameter to the nearest integer value, rounding according to the current rounding direction.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be rounded.

## Return Values

Upon successful completion, the **lrint**, **lrintf**, **lrintl**, **lrintd32**, **lrintd64**, and **lrintd128** subroutines return the rounded integer value.

If *x* is NaN, a domain error occurs and an unspecified value is returned.

If *x* is +Inf, a domain error occurs and an unspecified value is returned.

If *x* is -Inf, a domain error occurs and an unspecified value is returned.

If the correct value is positive and too large to represent as a long, a domain error occurs and an unspecified value is returned.

If the correct value is negative and too large to represent as a long, a domain error occurs and an unspecified value is returned.

## Related Information

“**feclearexcept** Subroutine” on page 265, “**fetestexcept** Subroutine” on page 273, and “**llrint**, **llrintf**, **llrintl**, **llrintd32**, **llrintd64**, and **llrintd128** Subroutines” on page 725.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## **lround, lroundf, lroundl, lroundd32, lroundd64, and lroundd128 Subroutines**

### **Purpose**

Rounds to the nearest integer value.

### **Syntax**

```
#include <math.h>
```

```
long lround (x)
double x;
```

```
long lroundf (x)
float x;
```

```
long lroundl (x)
long double x;
```

```
long lroundd32(x)
_Decimal32 x;
```

```
long lroundd64(x)
_Decimal64 x;
```

```
long lroundd128(x)
_Decimal128 x;
```

### **Description**

The **lround**, **lroundf**, **lroundl**, **lroundd32**, **lroundd64**, and **lroundd128** subroutines round the *x* parameter to the nearest integer value, rounding halfway cases away from zero, regardless of the current rounding direction.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### **Parameters**

*x* Specifies the value to be rounded.

### **Return Values**

Upon successful completion, the **lround**, **lroundf**, **lroundl**, **lroundd32**, **lroundd64**, and **lroundd128** subroutines return the rounded integer value.

If *x* is NaN, a domain error occurs and an unspecified value is returned.

If *x* is +Inf, a domain error occurs and an unspecified value is returned.

If *x* is -Inf, a domain error occurs and an unspecified value is returned.

If the correct value is positive and too large to represent as a **long**, a domain error occurs and an unspecified value is returned.

If the correct value is negative and too large to represent as a **long**, a domain error occurs and an unspecified value is returned.

## Related Information

“feclearexcept Subroutine” on page 265, “fetestexcept Subroutine” on page 273, and “llround, llroundf, llroundl, llroundd32, llroundd64, and llroundd128 Subroutines” on page 726.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## Isearch or lfind Subroutine

### Purpose

Performs a linear search and update.

### Library

Standard C Library (**libc.a**)

### Syntax

```
void *lsearch (Key, Base, NumberOfElementsPointer, Width, ComparisonPointer)
const void *Key;
void *Base;
size_t Width, *NumberOfElementsPointer;
int (*ComparisonPointer) (const void*, const void*);

void *lfind (Key, Base, NumberOfElementsPointer, Width, ComparisonPointer)
const void *Key, Base;
size_t Width, *NumberOfElementsPointer;
int (*ComparisonPointer) (const void*, const void*);
```

### Description

**Warning:** Undefined results can occur if there is not enough room in the table for the **lsearch** subroutine to add a new item.

The **lsearch** subroutine performs a linear search.

The algorithm returns a pointer to a table where data can be found. If the data is not in the table, the program adds it at the end of the table.

The **lfind** subroutine is identical to the **lsearch** subroutine, except that if the data is not found, it is not added to the table. In this case, a NULL pointer is returned.

The pointers to the *Key* parameter and the element at the base of the table should be of type pointer-to-element and cast to type pointer-to-character. The value returned should be cast into type pointer-to-element.

The comparison function need not compare every byte; therefore, the elements can contain arbitrary data in addition to the values being compared.

### Parameters

<i>Base</i>	Points to the first element in the table.
<i>ComparisonPointer</i>	Specifies the name (that you supply) of the comparison function ( <b>strcmp</b> , for example). It is called with two parameters that point to the elements being compared.
<i>Key</i>	Specifies the data to be sought in the table.

*NumberOfElementsPointer*

Points to an integer containing the current number of elements in the table. This integer is incremented if the data is added to the table.

*Width*

Specifies the size of an element in bytes.

The comparison function compares its parameters and returns a value as follows:

- If the first parameter equals the second parameter, the *ComparisonPointer* parameter returns a value of 0.
- If the first parameter does not equal the second parameter, the *ComparisonPointer* parameter returns a value of 1.

## Return Values

If the sought entry is found, both the **lsearch** and **lfind** subroutines return a pointer to it. Otherwise, the **lfind** subroutine returns a null pointer and the **lsearch** subroutine returns a pointer to the newly added element.

## Related Information

The **bsearch** (“bsearch Subroutine” on page 123) subroutine, **hsearch** (“hsearch, hcreate, or hdestroy Subroutine” on page 527) subroutine, **qsort** subroutine, **tsearch** subroutine.

Donald E. Knuth. *The Art of Computer Programming*, Volume 3, 6.1, Algorithm S. Reading, Massachusetts: Addison-Wesley, 1981.

Searching and Sorting Example Program and Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## Iseek, lseek or lseek64 Subroutine

### Purpose

Moves the read-write file pointer.

### Library

Standard C Library (**libc.a**)

### Syntax

```
off_t lseek ( FileDescriptor, Offset, Whence)
int FileDescriptor, Whence;
off_t Offset;
offset_t llseek (FileDescriptor, Offset, Whence)
int FileDescriptor, Whence;
offset_t Offset;
off64_t lseek64 (FileDescriptor, Offset, Whence)
int FileDescriptor, Whence;
off64_t Offset;
```

### Description

The **lseek**, **llseek**, and **lseek64** subroutines set the read-write file pointer for the open file specified by the *FileDescriptor* parameter. The **lseek** subroutine limits the *Offset* to **OFF\_MAX**.

In the large file enabled programming environment, **lseek** subroutine is redefined to **lseek64**.

If the *FileDescriptor* parameter refers to a shared memory object, the **lseek** subroutine fails with **EINVAL**.

## Parameters

<i>FileDescriptor</i>	Specifies a file descriptor obtained from a successful <b>open</b> or <b>fcntl</b> subroutine.
<i>Offset</i>	Specifies a value, in bytes, that is used in conjunction with the <i>Whence</i> parameter to set the file pointer. A negative value causes seeking in the reverse direction.
<i>Whence</i>	Specifies how to interpret the <i>Offset</i> parameter by setting the file pointer associated with the <i>FileDescriptor</i> parameter to one of the following variables: <b>SEEK_SET</b> Sets the file pointer to the value of the <i>Offset</i> parameter. <b>SEEK_CUR</b> Sets the file pointer to its current location plus the value of the <i>Offset</i> parameter. <b>SEEK_END</b> Sets the file pointer to the size of the file plus the value of the <i>Offset</i> parameter.

## Return Values

Upon successful completion, the resulting pointer location, measured in bytes from the beginning of the file, is returned. If either the **lseek** or **llseek** subroutines are unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **lseek** or **llseek** subroutines are unsuccessful and the file pointer remains unchanged if any of the following are true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not an open file descriptor.
<b>EINVAL</b>	The resulting offset would be greater than the maximum offset allowed for the file or device associated with <i>FileDescriptor</i> . The <b>lseek</b> subroutine was used with a file descriptor obtained from a call to the <b>shm_open</b> subroutine.
<b>EINVAL</b>	<i>Whence</i> is not one of the supported values.
<b>EOVERFLOW</b>	The resulting offset is larger than can be returned properly.
<b>ESPIPE</b>	The <i>FileDescriptor</i> parameter is associated with a pipe (FIFO) or a socket.

## Files

**/usr/include/unistd.h** Defines standard macros, data types and subroutines.

## Related Information

The **fcntl** (“fcntl, dup, or dup2 Subroutine” on page 255) subroutine, **fseek**, **rewind**, **ftell**, **fgetpos**, or **fsetpos** (“fseek, fseeko, fseeko64, rewind, ftell, ftello, ftello64, fgetpos, fgetpos64, fsetpos, or fsetpos64 Subroutine” on page 318) subroutine, **open**, **openx**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **read**, **readx**, **readv**, or **readvx** subroutine, **write**, **writex**, **writev**, or **writevx** subroutine.

File Systems and Directories in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## Ivm\_querylv Subroutine

### Purpose

Queries a logical volume and returns all pertinent information.

## Library

Logical Volume Manager Library (**liblvm.a**)

## Syntax

```
#include <lvm.h>
```

```
int lvm_querylv ( LV_ID, QueryLV, PVName)
struct lv_id *LV_ID;
struct querylv **QueryLV;
char *PVName;
```

## Description

**Note:** The **lvm\_querylv** subroutine uses the **sysconfig** system call, which requires root user authority, to query and update kernel data structures describing a volume group. You must have root user authority to use the **lvm\_querylv** subroutine.

The **lvm\_querylv** subroutine returns information for the logical volume specified by the *LV\_ID* parameter.

The **querylv** structure, found in the **lvm.h** file, is defined as follows:

```
struct querylv {
    char lvname[LVM_NAMESIZ];
    struct unique_id vg_id;
    long maxsize;
    long mirror_policy;
    long lv_state;
    long currentsize;
    long ppsize;
    long permissions;
    long bb_relocation;
    long write_verify;
    long mirwrt_consist;
    long open_close;
    struct pp *mirrors[LVM_NUMCOPIES];
    unsigned int stripe_exp;
    unsigned int striping_width;
}
struct pp {
    struct unique_id pv_id;
    long lp_num;
    long pp_num;
    long ppstate;
}
```

Field	Description
<b>lvname</b>	Specifies the special file name of the logical volume and can be either the full path name or a single file name that must reside in the <b>/dev</b> directory (for example, <b>rhdt1</b> ). All name fields must be null-terminated strings of from 1 to <b>LVM_NAMESIZ</b> bytes, including the null byte. If a raw or character device is not specified for the <b>lvname</b> field, the Logical Volume Manager (LVM) will add an <b>r</b> to the file name to have a raw device name. If there is no raw device entry for this name, the LVM will return the <b>LVM_NOTCHARDEV</b> error code.
<b>vg_id</b>	Specifies the unique ID of the volume group that contains the logical volume.
<b>maxsize</b>	Indicates the maximum size in logical partitions for the logical volume and must be in the range of 1 to <b>LVM_MAXLPS</b> .
<b>mirror_policy</b>	Specifies how the physical copies are written. The <b>mirror_policy</b> field should be either <b>LVM_SEQUENTIAL</b> or <b>LVM_PARALLEL</b> to indicate how the physical copies of a logical partition are to be written when there is more than one copy.

<b>Field</b>	<b>Description</b>
<b>lv_state</b>	<p>Specifies the current state of the logical volume and can have any of the following bit-specific values ORed together:</p> <p><b>LVM_LVDEFINED</b> The logical volume is defined.</p> <p><b>LVM_LVSTALE</b> The logical volume contains stale partitions.</p>
<b>currentsize</b>	<p>Indicates the current size in logical partitions of the logical volume. The size, in bytes, of every physical partition is 2 to the power of the <b>ppsize</b> field.</p>
<b>ppsize</b>	<p>Specifies the size of the physical partitions of all physical volumes in the volume group.</p>
<b>permissions</b>	<p>Specifies the permission assigned to the logical volume and can be one of the following values:</p> <p><b>LVM_RDONLY</b> Access to this logical volume is read only.</p> <p><b>LVM_RDWR</b> Access to this logical volume is read/write.</p>
<b>bb_relocation</b>	<p>Specifies if bad block relocation is desired and is one of the following values:</p> <p><b>LVM_NORELOC</b> Bad blocks will not be relocated.</p> <p><b>LVM_RELOC</b> Bad blocks will be relocated.</p>
<b>write_verify</b>	<p>Specifies if write verification for the logical volume is desired and returns one of the following values:</p> <p><b>LVM_NOVERIFY</b> Write verification is not performed for this logical volume.</p> <p><b>LVM_VERIFY</b> Write verification is performed on all writes to the logical volume.</p>
<b>mirwrt_consist</b>	<p>Indicates whether mirror-write consistency recovery will be performed for this logical volume.</p> <p>The LVM always ensures data consistency among mirrored copies of a logical volume during normal I/O processing. For every write to a logical volume, the LVM generates a write request for every mirror copy. A problem arises if the system crashes in the middle of processing a mirrored write (before all copies are written). If mirror write consistency recovery is requested for a logical volume, the LVM keeps additional information to allow recovery of these inconsistent mirrors. Mirror write consistency recovery should be performed for most mirrored logical volumes. Logical volumes, such as page space, that do not use the existing data when the volume group is re-varied on do not need this protection.</p> <p>Values for the <b>mirwrt_consist</b> field are:</p> <p><b>LVM_CONSIST</b> Mirror-write consistency recovery will be done for this logical volume.</p> <p><b>LVM_NOCONSIST</b> Mirror-write consistency recovery will not be done for this logical volume.</p>
<b>open_close</b>	<p>Specifies if the logical volume is opened or closed. Values for this field are:</p> <p><b>LVM_QLV_NOTOPEN</b> The logical volume is closed.</p> <p><b>LVM_QLVOPEN</b> The logical volume is opened by one or more processes.</p>



Field	Description
<b>mirrors</b>	Specifies an array of pointers to partition map lists (physical volume id, logical partition number, physical partition number, and physical partition state for each copy of the logical partitions for the logical volume). The <b>ppstate</b> field can be <b>LVM_PPFREE</b> , <b>LVM_PPALLOC</b> , or <b>LVM_PPSTALE</b> . If a logical partition does not contain any copies, its <b>pv_id</b> , <b>lp_num</b> , and <b>pp_num</b> fields will contain zeros.
<b>stripe_exp</b>	Specifies the log base 2 of the logical volume strip size (the strip size multiplied by the number of disks in an array equals the stripe size). For example, $2^{20}$ is 1048576 (that is, 1 MB). Therefore, if the strip size is 1 MB, the <b>stripe_exp</b> field is 20. If the logical volume is not striped, the <b>stripe_exp</b> field is 0.
<b>stripe_width</b>	Specifies the number of disks that form the striped logical volume. If the logical volume is not striped, the <b>striping_width</b> field is 0.

The *PVName* parameter enables the user to query from a volume group descriptor area on a specific physical volume instead of from the Logical Volume Manager's (LVM) most recent, in-memory copy of the descriptor area. This method should only be used if the volume group is varied off.

**Note:** The data returned is not guaranteed to be the most recent or correct, and it can reflect a back-level descriptor area.

The *PVName* parameter should specify either the full path name of the physical volume that contains the descriptor area to query, or a single file name that must reside in the */dev* directory (for example, **rhdisk1**). This parameter must be a null-terminated string between 1 and **LVM\_NAMESIZ** bytes, including the null byte, and must represent a raw device entry. If a raw or character device is not specified for the *PVName* parameter, the LVM adds an **r** to the file name to have a raw device name. If there is no raw device entry for this name, the LVM returns the **LVM\_NOTCHARDEV** error code.

If a *PVName* parameter is specified, only the **minor\_num** field of the *LV\_ID* parameter need be supplied. The LVM fills in the **vg\_id** field and returns it to the user. If the user wishes to query from the LVM's in-memory copy, the *PVName* parameter should be set to null. When using this method of query, the volume group must be varied on, or an error is returned.

**Note:** As long as the *PVName* parameter is not null, the LVM will attempt a query from a physical volume and *not* from its in-memory copy of data.

In addition to the *PVName* parameter, the caller passes the ID of the logical volume to be queried (*LV\_ID* parameter) and the address of a pointer to the **querylv** structure, specified by the *QueryLV* parameter. The LVM separately allocates the space needed for the **querylv** structure and the struct **pp** arrays, and returns the **querylv** structure's address in the pointer variable passed in by the user. The user is responsible for freeing the space by first freeing the struct **pp** pointers in the **mirrors** array and then freeing the **querylv** structure.

**Attention:** To prevent corruption when there are many **pp** arrays, the caller of **lvm\_querylv** must set *QueryLV->mirrors k != NULL*.

## Parameters

<i>LV_ID</i>	Points to an <b>lv_id</b> structure that specifies the logical volume to query.
<i>QueryLV</i>	Contains the address of a pointer to the <b>querylv</b> structure.
<i>PVName</i>	Names the physical volume from which to use the volume group descriptor for the query. This parameter can also be null.

## Return Values

If the **lvm\_querylv** subroutine is successful, it returns a value of 0.

## Error Codes

If the `lvm_querylv` subroutine does not complete successfully, it returns one of the following values:

<code>LVM_ALLOCERR</code>	The subroutine could not allocate enough space for the complete buffer.
<code>LVM_INVALID_MIN_NUM</code>	The minor number of the logical volume is not valid.
<code>LVM_INVALID_PARAM</code>	A parameter passed into the routine is not valid.
<code>LVM_INV_DEVENT</code>	The device entry for the physical volume specified by the <code>Pvname</code> parameter is not valid and cannot be checked to determine if it is raw.
<code>LVM_NOTCHARDEV</code>	The physical volume name given does not represent a raw or character device.
<code>LVM_OFFLINE</code>	The volume group containing the logical volume to query was offline.
	If the query originates from the varied-on volume group's current volume group descriptor area, one of the following error codes is returned:
<code>LVM_DALVOPN</code>	The volume group reserved logical volume could not be opened.
<code>LVM_MAPFBSY</code>	The volume group is currently locked because system management on the volume group is being done by another process.
<code>LVM_MAPFOPN</code>	The mapped file, which contains a copy of the volume group descriptor area used for making changes to the volume group, could not be opened.
<code>LVM_MAPFRDWR</code>	The mapped file could not be read or written.

If a physical volume name has been passed, requesting that the query originate from a specific physical volume, one of the following error codes is returned:

<code>LVM_BADBBDIR</code>	The bad-block directory could not be read or written.
<code>LVM_LVMRECERR</code>	The LVM record, which contains information about the volume group descriptor area, could not be read.
<code>LVM_NOPVVGDA</code>	There are no volume group descriptor areas on the physical volume specified.
<code>LVM_NOTVGMEM</code>	The physical volume specified is not a member of a volume group.
<code>LVM_PVDAREAD</code>	An error occurred while trying to read the volume group descriptor area from the specified physical volume.
<code>LVM_PVOPNERR</code>	The physical volume device could not be opened.
<code>LVM_VGDA_BB</code>	A bad block was found in the volume group descriptor area located on the physical volume that was specified for the query. Therefore, a query cannot be done from the specified physical volume.

## Related Information

List of Logical Volume Subroutines and Logical Volume Programming Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## `lvm_querypv` Subroutine

### Purpose

Queries a physical volume and returns all pertinent information.

### Library

Logical Volume Manager Library (`liblvm.a`)

### Syntax

```
#include <lvm.h>
```

```
int lvm_querypv (VG_ID, PV_ID, QueryPV, PVName)
struct unique_id * VG_ID;
```

```

struct unique_id * PV_ID;
struct querypv ** QueryPV;
char * PVName;

```

## Description

**Note:** The `lvm_querypv` subroutine uses the `sysconfig` system call, which requires root user authority, to query and update kernel data structures describing a volume group. You must have root user authority to use the `lvm_querypv` subroutine.

The `lvm_querypv` subroutine returns information on the physical volume specified by the `PV_ID` parameter.

The `querypv` structure, defined in the `lvm.h` file, contains the following fields:

```

struct querypv {
    long ppsize;
    long pv_state;
    long pp_count;
    long alloc_ppcount;
    long pvnum_vgdas;
    struct pp_map *pp_map;
    char hotspare;
    struct unique_id pv_id;
    long freespace;
}
struct pp_map {
    long pp_state;
    struct lv_id lv_id;
    long lp_num;
    long copy;
    struct unique_id fst_alt_vol;
    long fst_alt_part;
    struct unique_id snd_alt_vol;
    long snd_alt_part;
}

```

Field	Description
ppsize	Specifies the size of the physical partitions, which is the same for all partitions within a volume group. The size in bytes of a physical partition is 2 to the power of ppsize.
pv_state	Contains the current state of the physical volume.
pp_count	Contains the total number of physical partitions on the physical volume.
alloc_ppcount	Contains the number of allocated physical partitions on the physical volume.

<b>Field</b>	<b>Description</b>
pp_map	<p>Points to an array that has entries for each physical partition of the physical volume. Each entry in this array will contain the pp_state that specifies the state of the physical partition (<b>LVM_PPFREE</b>, <b>LVM_PPALLOC</b>, or <b>LVM_PPSTALE</b>) and the lv_id, field, the ID of the logical volume that it is a member of. The pp_map array also contains the physical volume IDs (fst_alt_vol and snd_alt_vol) and the physical partition numbers (fst_alt_part and snd_alt_part) for the first and second alternate copies of the physical partition, and the logical partition number (lp_num) that the physical partition corresponds to.</p> <p>If the physical partition is free (that is, not allocated), <i>all</i> of its <b>pp_map</b> fields will be zero.</p> <p><b>fst_alt_vol</b> Contains zeros if the logical partition has only one physical copy.</p> <p><b>fst_alt_part</b> Contains zeros if the logical partition has only one physical copy.</p> <p><b>snd_alt_vol</b> Contains zeros if the logical partition has only one or two physical copies.</p> <p><b>snd_alt_part</b> Contains zeros if the logical partition has only one or two physical copies.</p> <p><b>copy</b> Specifies which copy of a logical partition this physical partition is allocated to. This field will contain one of the following values:</p> <p><b>LVM_PRIMARY</b> Primary and only copy of a logical partition</p> <p><b>LVM_PRIMOF2</b> Primary copy of a logical partition with two physical copies</p> <p><b>LVM_PRIMOF3</b> Primary copy of a logical partition with three physical copies</p> <p><b>LVM_SCNDOF2</b> Secondary copy of a logical partition with two physical copies</p> <p><b>LVM_SCNDOF3</b> Secondary copy of a logical partition with three physical copies</p> <p><b>LVM_TERTOF3</b> Tertiary copy of a logical partition with three physical copies.</p>
pvnum_vgdas	Contains the number of volume group descriptor areas (0, 1, or 2) that are on the specified physical volume.
hotspare	Specifies that the physical volume is a hotspare.
pv_id	Specifies the physical volume identifier.
freespace	Specifies the number of physical partitions in the volume group.

The *PVName* parameter enables the user to query from a volume group descriptor area on a specific physical volume instead of from the Logical Volume Manager's (LVM) most recent, in-memory copy of the descriptor area. This method should only be used if the volume group is varied off. The data returned is not guaranteed to be most recent or correct, and it can reflect a back level descriptor area.

The *PVname* parameter should specify either the full path name of the physical volume that contains the descriptor area to query or a single file name that must reside in the */dev* directory (for example, **rhdisk1**). This field must be a null-terminated string of from 1 to **LVM\_NAMESIZ** bytes, including the null byte, and represent a raw or character device. If a raw or character device is not specified for the *PVName* parameter, the LVM will add an **r** to the file name in order to have a raw device name. If there is no raw device entry for this name, the LVM will return the **LVM\_NOTCHARDEV** error code. If a *PVName* is specified, the volume group identifier, VG\_ID, will be returned by the LVM through the *VG\_ID* parameter

passed in by the user. If the user wishes to query from the LVM in-memory copy, the *PVName* parameter should be set to null. When using this method of query, the volume group must be varied on, or an error will be returned.

**Note:** As long as the *PVName* is not null, the LVM will attempt a query from a physical volume and *not* from its in-memory copy of data.

In addition to the *PVName* parameter, the caller passes the *VG\_ID* parameter, indicating the volume group that contains the physical volume to be queried, the unique ID of the physical volume to be queried, the *PV\_ID* parameter, and the address of a pointer of the type *QueryPV*. The LVM will separately allocate enough space for the **querypv** structure and the struct *pp\_map* array and return the address of the **querypv** structure in the *QueryPV* pointer passed in. The user is responsible for freeing the space by freeing the struct *pp\_map* pointer and then freeing the *QueryPV* pointer.

## Parameters

<i>VG_ID</i>	Points to a <b>unique_id</b> structure that specifies the volume group of which the physical volume to query is a member.
<i>PV_ID</i>	Points to a <b>unique_id</b> structure that specifies the physical volume to query.
<i>QueryPV</i>	Specifies the address of a pointer to a <b>querypv</b> structure.
<i>PVName</i>	Names a physical volume from which to use the volume group descriptor area for the query. This parameter can be null.

## Return Values

The **lvm\_querypv** subroutine returns a value of 0 upon successful completion.

## Error Codes

If the **lvm\_querypv** subroutine fails it returns one of the following error codes:

<b>LVM_ALLOCERR</b>	The routine cannot allocate enough space for a complete buffer.
<b>LVM_INVALID_PARAM</b>	An invalid parameter was passed into the routine.
<b>LVM_INV_DEVENT</b>	The device entry for the physical volume is invalid and cannot be checked to determine if it is raw.
<b>LVM_OFFLINE</b>	The volume group specified is offline and should be online.

If the query originates from the varied-on volume group's current volume group descriptor area, one of the following error codes may be returned:

<b>LVM_DALVOPN</b>	The volume group reserved logical volume could not be opened.
<b>LVM_MAPFBSY</b>	The volume group is currently locked because system management on the volume group is being done by another process.
<b>LVM_MAPFOPN</b>	The mapped file, which contains a copy of the volume group descriptor area used for making changes to the volume group, could not be opened.
<b>LVM_MAPFRDWR</b>	Either the mapped file could not be read, or it could not be written.

If a physical volume name has been passed, requesting that the query originate from a specific physical volume, then one of the following error codes may be returned:

<b>LVM_BADBBDIR</b>	The bad-block directory could not be read or written.
<b>LVM_LVMRECERR</b>	The LVM record, which contains information about the volume group descriptor area, could not be read.
<b>LVM_NOPVVGDA</b>	There are no volume group descriptor areas on this physical volume.
<b>LVM_NOTCHARDEV</b>	A device is not a raw or character device.

<b>LVM_NOTVGMEM</b>	The physical volume is not a member of a volume group.
<b>LVM_PVDAREAD</b>	An error occurred while trying to read the volume group descriptor area from the specified physical volume.
<b>LVM_PVOPNERR</b>	The physical volume device could not be opened.
<b>LVM_VGDA_BB</b>	A bad block was found in the volume group descriptor area located on the physical volume that was specified for the query. Therefore, a query cannot be done from the specified physical volume.

## Related Information

List of Logical Volume Subroutines and Logical Volume Programming Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## lvm\_queryvg Subroutine

### Purpose

Queries a volume group and returns pertinent information.

### Library

Logical Volume Manager Library (**liblvm.a**)

### Syntax

```
#include <lvm.h>
```

```
int lvm_queryvg ( VG_ID, QueryVG, PVName)
struct unique_id *VG_ID;
struct queryvg **QueryVG;
char *PVName;
```

### Description

**Note:** The **lvm\_queryvg** subroutine uses the **sysconfig** system call, which requires root user authority, to query and update kernel data structures describing a volume group. You must have root user authority to use the **lvm\_queryvg** subroutine.

The **lvm\_queryvg** subroutine returns information on the volume group specified by the **VG\_ID** parameter.

The **queryvg** structure, found in the **lvm.h** file, contains the following fields:

```
struct queryvg {
    long maxlvs;
    long ppsize;
    long freespace;
    long num_lvs;
    long num_pvs;
    long total_vgdas;
    struct lv_array *lvs;
    struct pv_array *pvs;
    short conc_capable;
    short default_mode;
    short conc_status;
    unsigned int maxpvs;
    unsigned int maxpvpps;
    unsigned int maxvgpps;
}
struct pv_array {
    struct unique_id pv_id;
```

```

    char  state;
    char  res[3];
    long  pvnum_vgdas;
}
struct lv_array {
    struct lv_id  lv_id;
    char  lvname[LVM_NAMESIZ];
    char  state;
    char  res[3];
}

```

Field	Description
maxlvs	Specifies the maximum number of logical volumes allowed in the volume group.
ppsize	Specifies the size of all physical partitions in the volume group. The size in bytes of each physical partitions is 2 to the power of the ppsize field.
freespace	Contains the number of free physical partitions in this volume group.
num_lvs	Indicates the number of logical volumes.
num_pvs	Indicates the number of physical volumes.
total_vgdas	Specifies the total number of volume group descriptor areas for the entire volume group.
lvs	Points to an array of unique IDs, names, and states of the logical volumes in the volume group.
pvs	Points to an array of unique IDs, states, and the number of volume group descriptor areas for each of the physical volumes in the volume group.
conc_capable	Indicates that the volume group was created concurrent mode capable if the value is equal to 1.
default_mode	The behavior of this value is undefined.
conc_status	Indicates that the volume group is varied on in concurrent mode.
maxpvs	Specifies the maximum number of physical volumes allowed in the volume group.
maxpvpps	Specifies the maximum number of physical partitions allowed for a physical volume in the volume group.
maxvgpps	Specifies the maximum number of physical partitions allowed for the entire volume group.

The *PVName* parameter enables the user to query from a descriptor area on a specific physical volume instead of from the Logical Volume Manager's (LVM) most recent, in-memory copy of the descriptor area. This method should only be used if the volume group is varied off. The data returned is *not guaranteed* to be most recent or correct, and it can reflect a back level descriptor area. The *Pvname* parameter should specify either the full path name of the physical volume that contains the descriptor area to query or a single file name that must reside in the */dev* directory (for example, **rhdisk1**). The name must represent a raw device. If a raw or character device is not specified for the *PVName* parameter, the Logical Volume Manager will add an *r* to the file name in order to have a raw device name. If there is no raw device entry for this name, the LVM returns the **LVM\_NOTCHARDEV** error code. This field must be a null-terminated string of from 1 to **LVM\_NAMESIZ** bytes, including the null byte. If a *PVName* is specified, the LVM will return the *VG\_ID* to the user through the *VG\_ID* pointer passed in. If the user wishes to query from the LVM in-memory copy, the *PVName* parameter should be set to null. When using this method of query, the volume group must be varied on, or an error will be returned.

**Note:** As long as the *PVName* parameter is not null, the LVM will attempt a query from a physical volume and *not* its in-memory copy of data.

In addition to the *PVName* parameter, the caller passes the unique ID of the volume group to be queried (*VG\_ID*) and the address of a pointer to a **queryvg** structure. The LVM will separately allocate enough space for the **queryvg** structure, as well as the **lv\_array** and **pv\_array** structures, and return the address of the completed structure in the *QueryVG* parameter passed in by the user. The user is responsible for freeing the space by freeing the *lv* and *pv* pointers and then freeing the *QueryVG* pointer.

## Parameters

<i>VG_ID</i>	Points to a <b>unique_id</b> structure that specifies the volume group to be queried.
<i>QueryVG</i>	Specifies the address of a pointer to the <b>queryvg</b> structure.
<i>PVName</i>	Specifies the name of the physical volume that contains the descriptor area to query and must be the name of a raw device.

## Return Values

The **lvm\_queryvg** subroutine returns a value of 0 upon successful completion.

## Error Codes

If the **lvm\_queryvg** subroutine fails it returns one of the following error codes:

<b>LVM_ALLOCERR</b>	The subroutine cannot allocate enough space for a complete buffer.
<b>LVM_FORCEOFF</b>	The volume group has been forcefully varied off due to a loss of quorum.
<b>LVM_INVALID_PARAM</b>	An invalid parameter was passed into the routine.
<b>LVM_OFFLINE</b>	The volume group is offline and should be online.

If the query originates from the varied-on volume group's current volume group descriptor area, one of the following error codes may be returned:

<b>LVM_DALVOPN</b>	The volume group reserved logical volume could not be opened.
<b>LVM_INV_DEVENT</b>	The device entry for the physical volume specified by the <i>PVName</i> parameter is invalid and cannot be checked to determine if it is raw.
<b>LVM_MAPFBSY</b>	The volume group is currently locked because system management on the volume group is being done by another process.
<b>LVM_MAPFOPN</b>	The mapped file, which contains a copy of the volume group descriptor area used for making changes to the volume group, could not be opened.
<b>LVM_MAPFRDWR</b>	Either the mapped file could not be read, or it could not be written.
<b>LVM_NOTCHARDEV</b>	A device is not a raw or character device.

If a physical volume name has been passed, requesting that the query originate from a specific physical volume, one of the following error codes may be returned:

<b>LVM_BADBBDIR</b>	The bad-block directory could not be read or written.
<b>LVM_LVMRECERR</b>	The LVM record, which contains information about the volume group descriptor area, could not be read.
<b>LVM_NOPVVGDA</b>	There are no volume group descriptor areas on this physical volume.
<b>LVM_NOTVGMEM</b>	The physical volume is not a member of a volume group.
<b>LVM_PVDAREAD</b>	An error occurred while trying to read the volume group descriptor area from the specified physical volume.
<b>LVM_PVOPNERR</b>	The physical volume device could not be opened.
<b>LVM_VGDA_BB</b>	A bad block was found in the volume group descriptor area located on the physical volume that was specified for the query. Therefore, a query cannot be done from this physical volume.

## Related Information

List of Logical Volume Subroutines and Logical Volume Programming Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.



---

## Ivm\_queryvgs Subroutine

### Purpose

Queries volume groups and returns information to online volume groups.

### Library

Logical Volume Manager Library (**liblvm.a**)

### Syntax

```
#include <lvm.h>
```

```
int lvm_queryvgs ( QueryVGS, Kmid)
struct queryvgs **QueryVGS;
mid_t Kmid;
```

### Description

**Note:** The **lvm\_queryvgs** subroutine uses the **sysconfig** system call, which requires root user authority, to query and update kernel data structures describing a volume group. You must have root user authority to use the **lvm\_queryvgs** subroutine.

The **lvm\_queryvgs** subroutine returns the volume group IDs and major numbers for all volume groups in the system that are online.

The caller passes the address of a pointer to a **queryvgs** structure, and the Logical Volume Manager (LVM) allocates enough space for the structure and returns the address of the structure in the pointer passed in by the user. The caller also passes in a *Kmid* parameter, which identifies the entry point of the logical device driver module:

```
struct queryvgs {
    long num_vgs;
    struct {
        long major_num
        struct unique_id vg_id;
    } vgs [LVM_MAXVGS];
}
```

Field	Description
num_vgs	Contains the number of online volume groups on the system. The vgs is an array of the volume group IDs and major numbers of all online volume groups in the system.

### Parameters

<i>QueryVGS</i>	Points to the <b>queryvgs</b> structure.
<i>Kmid</i>	Identifies the address of the entry point of the logical volume device driver module.

### Return Values

The **lvm\_queryvgs** subroutine returns a value of 0 upon successful completion.

### Error Codes

If the **lvm\_queryvgs** subroutine fails, it returns one of the following error codes:

<b>LVM_ALLOCERR</b>	The routine cannot allocate enough space for the complete buffer.
---------------------	---

LVM\_INVALID\_PARAM  
LVM\_INVCONFIG

An invalid parameter was passed into the routine.

An error occurred while attempting to configure this volume group into the kernel. This error will normally result if the module ID is invalid, if the major number given is already in use, or if the volume group device could not be opened.

## Related Information

List of Logical Volume Subroutines and Logical Volume Programming Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## malloc, free, realloc, calloc, malloc, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine

### Purpose

Provides a complete set of memory allocation, reallocation, deallocation, and heap management tools.

### Libraries

Berkeley Compatibility Library (**libbsd.a**)

Standard C Library (**libc.a**)

### Malloc Subsystem APIs

- malloc
- free
- realloc
- calloc
- malloc
- mallinfo
- mallinfo\_heap
- alloca
- valloc
- posix\_memalign

### malloc

### Syntax

```
#include <stdlib.h>
```

```
void *malloc (Size)  
size_t Size;
```

### Description

The **malloc** subroutine returns a pointer to a block of memory of at least the number of bytes specified by the *Size* parameter. The block is aligned so that it can be used for any type of data. Undefined results occur if the space assigned by the **malloc** subroutine is overrun.

### Parameters

*Size* Specifies the size, in bytes, of memory to allocate.

## Return Values

Upon successful completion, the **malloc** subroutine returns a pointer to space suitably aligned for the storage of any type of object. If the size requested is 0, **malloc** returns NULL in normal circumstances. However, if the program was compiled with the defined **\_LINUX\_SOURCE\_COMPAT** macro, **malloc** returns a valid pointer to a space of size 0.

If the request cannot be satisfied for any reason, the **malloc** subroutine returns NULL.

## Error Codes

**ENOMEM**                      Insufficient storage space is available to service the request.

## free

### Syntax

```
#include <stdlib.h>

void free (Pointer)
void * Pointer;
```

### Description

The **free** subroutine deallocates a block of memory previously allocated by the **malloc** subsystem. Undefined results occur if the *Pointer* parameter is not an address that has previously been allocated by the **malloc** subsystem, or if the *Pointer* parameter has already been deallocated. If the *Pointer* parameter is NULL, no action occurs.

### Parameters

*Pointer*                                      Specifies a pointer to space previously allocated by the **malloc** subsystem.

## Return Values

The **free** subroutine does not return a value. Upon successful completion with nonzero arguments, the **realloc** subroutine returns a pointer to the (possibly moved) allocated space. If the *Size* parameter is 0 and the *Pointer* parameter is not null, no action occurs.

## Error Codes

The **free** subroutine does not set **errno**.

## realloc

### Syntax

```
#include <stdlib.h>

void *realloc (Pointer, Size)
void *Pointer;
size_t Size;
```

## Description

The **realloc** subroutine changes the size of the memory object pointed to by the *Pointer* parameter to the number of bytes specified by the *Size* parameter. The *Pointer* must point to an address returned by a **malloc** subsystem allocation routine, and must not have been previously deallocated. Undefined results occur if *Pointer* does not meet these criteria.

The contents of the memory object remain unchanged up to the lesser of the old and new sizes. If the current memory object cannot be enlarged to satisfy the request, the **realloc** subroutine acquires a new memory object and copies the existing data to the new space. The old memory object is then freed. If no memory object can be acquired to accommodate the request, the object remains unchanged.

If the *Pointer* parameter is null, the **realloc** subroutine is equivalent to a **malloc** subroutine of the same size.

If the *Size* parameter is 0 and the *Pointer* parameter is not null, the **realloc** subroutine is equivalent to a **free** subroutine of the same size.

## Parameters

<i>Pointer</i>	Specifies a <i>Pointer</i> to space previously allocated by the <b>malloc</b> subsystem.
<i>Size</i>	Specifies the new size, in bytes, of the memory object.

## Return Values

Upon successful completion with nonzero arguments, the **realloc** subroutine returns a pointer to the (possibly moved) allocated space. If the *Size* parameter is 0 and the *Pointer* parameter is not null, return behavior is equivalent to that of the **free** subroutine. If the *Pointer* parameter is null and the *Size* parameter is not zero, return behavior is equivalent to that of the **malloc** subroutine.

## Error Codes

ENOMEM	Insufficient storage space is available to service the request.
--------	---

## calloc

### Syntax

```
#include <stdlib.h>
```

```
void *calloc (NumberOfElements, ElementSize)  
size_t NumberOfElements;  
size_t ElementSize;
```

### Description

The **calloc** subroutine allocates space for an array containing the *NumberOfElements* objects. The *ElementSize* parameter specifies the size of each element in bytes. After the array is allocated, all bits are initialized to 0.

The order and contiguity of storage allocated by successive calls to the **calloc** subroutine is unspecified. The pointer returned points to the first (lowest) byte address of the allocated space. The allocated space is aligned so that it can be used for any type of data. Undefined results occur if the space assigned by the **calloc** subroutine is overrun.

## Parameters

<i>NumberOfElements</i>	Specifies the number of elements in the array.
<i>ElementSize</i>	Specifies the size, in bytes, of each element in the array.

## Return Values

Upon successful completion, the **calloc** subroutine returns a pointer to the allocated, zero-initialized array. If the size requested is 0, the **calloc** subroutine returns NULL in normal circumstances. However, if the program was compiled with the macro **\_LINUX\_SOURCE\_COMPAT** defined, the **calloc** subroutine returns a valid pointer to a space of size 0.

If the request cannot be satisfied for any reason, the **calloc** subroutine returns NULL.

## Error Codes

<b>ENOMEM</b>	Insufficient storage space is available to service the request.
---------------	---

## mallopt

### Syntax

```
#include <malloc.h>
#include <stdlib.h>
```

```
int mallopt (Command, Value)
int Command;
int Value;
```

### Description

The **mallopt** subroutine is provided for source-level compatibility with the System V **malloc** subroutine. The **mallopt** subroutine supports the following commands:

Command	Value	Effect
<b>M_MXFAST</b>	0	If called before any other <b>malloc</b> subsystem subroutine, this enables the Default allocation policy for the process.
<b>M_MXFAST</b>	1	If called before any other <b>malloc</b> subsystem subroutine, this enables the 3.1 allocation policy for the process.
<b>M_DISCLAIM</b>	0	If called while the Default Allocator is enabled, all free memory in the process heap is disclaimed.
<b>M_MALIGN</b>	<i>N</i>	If called at runtime, sets the default <b>malloc</b> allocation alignment to the value <i>N</i> . The <i>N</i> value must be a power of 2 (greater than or equal to the size of a pointer).

## Parameters

<i>Command</i>	Specifies the <b>mallopt</b> command to be executed.
<i>Value</i>	Specifies the size of each element in the array.

## Return Values

Upon successful completion, the **mallopt** subroutine returns 0. Otherwise, 1 is returned. If an invalid alignment is requested (one that is not a power of 2), **mallopt** fails with a return value of 1, although subsequent calls to **malloc** are unaffected and continue to provide the alignment value from before the failed **mallopt** call.

## Error Codes

The **mallopt** subroutine does not set **errno**.

## mallinfo

### Syntax

```
#include <malloc.h>
#include <stdlib.h>
```

```
struct mallinfo mallinfo();
```

### Description

The **mallinfo** subroutine can be used to obtain information about the heap managed by the **malloc** subsystem.

## Return Values

The **mallinfo** subroutine returns a structure of type **struct mallinfo**, filled in with relevant information and statistics about the heap. The contents of this structure can be interpreted using the definition of **struct mallinfo** in the `/usr/include/malloc.h` file.

## Error Codes

The **mallinfo** subroutine does not set **errno**.

## mallinfo\_heap

### Syntax

```
#include <malloc.h>
#include <stdlib.h>
```

```
struct mallinfo_heap mallinfo_heap (Heap)
int Heap;
```

### Description

In a multiheap context, the **mallinfo\_heap** subroutine can be used to obtain information about a specific heap managed by the **malloc** subsystem.

## Parameters

*Heap* Specifies which heap to query.

## Return Values

**mallinfo\_heap** returns a structure of type **struct mallinfo\_heap**, filled in with relevant information and statistics about the heap. The contents of this structure can be interpreted using the definition of **struct mallinfo\_heap** in the `/usr/include/malloc.h` file.

## Error Codes

The `mallinfo_heap` subroutine does not set `errno`.

## alloca

### Syntax

```
#include <stdlib.h>
```

```
char *alloca (Size)  
int Size;
```

### Description

The `alloca` subroutine returns a pointer to a block of memory of at least the number of bytes specified by the `Size` parameter. The space is allocated from the stack frame of the caller and is automatically freed when the calling subroutine returns.

If `alloca` is used in code compiled with the C++ compiler, `#pragma alloca` must be added to the source before the reference to `alloca`. Alternatively, the `-ma` compiler option can be used during compilation.

### Parameters

*Size* Specifies the size, in bytes, of memory to allocate.

### Return Values

The `alloca` subroutine returns a pointer to space of the requested size.

### Error Codes

The `alloca` subroutine does not set `errno`.

## valloc

### Syntax

```
#include <stdlib.h>
```

```
void *valloc (Size)  
size_t Size;
```

### Description

The `valloc` subroutine is supported as a compatibility interface in the Berkeley Compatibility Library (`libbsd.a`), as well as in `libc.a`. The `valloc` subroutine has the same effect as `malloc`, except that the allocated memory is aligned to a multiple of the value returned by `sysconf (_SC_PAGESIZE)`.

### Parameters

*Size* Specifies the size, in bytes, of memory to allocate.

## Return Values

Upon successful completion, the **valloc** subroutine returns a pointer to a memory object that is *Size* bytes in length, aligned to a page-boundary. Undefined results occur if the space assigned by the **valloc** subroutine is overrun.

If the request cannot be satisfied for any reason, **valloc** returns NULL.

## Error Codes

**ENOMEM**                      Insufficient storage space is available to service the request.

## posix\_memalign

### Syntax

```
#include <stdlib.h>
```

```
int posix_memalign(void **Pointer2Pointer, Align, Size)
void ** Pointer2Pointer;
size_t Align;
size_t Size;
```

### Description

The **posix\_memalign** subroutine allocates *Size* bytes of memory aligned on a boundary specified by *Align*. The address of this memory is stored in *Pointer2Pointer*.

### Parameters

<i>Pointer2Pointer</i>	Specifies the location in which the address should be copied.
<i>Align</i>	Specifies the alignment of the allocated memory, in bytes. The <i>Align</i> parameter must be a power-of-two multiple of the size of a pointer.
<i>Size</i>	Specifies the size, in bytes, of memory to allocate.

## Return Values

Upon successful completion, **posix\_memalign** returns 0. Otherwise, an error number is returned to indicate the error.

## Error Codes

**EINVAL**                      The value of *Align* is not a power-of-two multiple of the size of a pointer.  
**ENOMEM**                      Insufficient storage space is available to service the request.

## Related Information

The **\_end**, **\_etext**, or **\_edata** (“\_end, \_etext, or \_edata Identifier” on page 224) identifier.

User Defined Malloc Replacement, Debug Malloc, Malloc Multiheap, Malloc Buckets, Malloc Log, Malloc Trace, System Memory Allocation Using the malloc Subsystem, Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*, and Paging space and virtual memory in *Operating system and device management*.



---

# madd, msub, mult, mdiv, pow, gcd, invert, rpow, msqrt, mcmp, move, min, omin, fmin, m\_in, mout, omout, fmout, m\_out, sdiv, or itom Subroutine

## Purpose

Multiple-precision integer arithmetic.

## Library

Berkeley Compatibility Library (**libbsd.a**)

## Syntax

```
#include <mp.h>
#include <stdio.h>
```

```
typedef struct mint {int Length; short * Value} MINT;
```

```
madd( a, b, c)
msub(a,b,c)
mult(a,b,c)
mdiv(a,b, q, r)
pow(a,b, m,c)
gcd(a,b,c)
invert(a,b,c)
rpow(a,n,c)
msqrt(a,b,r)
mcmp(a,b)
move(a,b)
min(a)
omin(a)
fmin(a,f)
m_in(a, n,f)
mout(a)
omout(a)
fmout(a,f)
m_out(a,n,f)
MINT *a, *b, *c, *m, *q, *r;
FILE * f;
int n;
sdiv(a,n,q,r)
MINT *a, *q;
short n;
short *r;
MINT *itom(n)
```

## Description

These subroutines perform arithmetic on integers of arbitrary *Length*. The integers are stored using the defined type **MINT**. Pointers to a **MINT** can be initialized using the **itom** subroutine, which sets the initial *Value* to *n*. After that, space is managed automatically by the subroutines.

The **madd** subroutine, **msub** subroutine, and **mult** subroutine assign to *c* the sum, difference, and product, respectively, of *a* and *b*.

The **mdiv** subroutine assigns to *q* and *r* the quotient and remainder obtained from dividing *a* by *b*.

The **sdiv** subroutine is like the **mdiv** subroutine except that the divisor is a short integer  $n$  and the remainder is placed in a short whose address is given as  $r$ .

The **msqrt** subroutine produces the integer square root of  $a$  in  $b$  and places the remainder in  $r$ .

The **rpow** subroutine calculates in  $c$  the value of  $a$  raised to the (regular integral) power  $n$ , while the **pow** subroutine calculates this with a full multiple precision exponent  $b$  and the result is reduced modulo  $m$ .

**Note:** The **pow** subroutine is also present in the IEEE Math Library, **libm.a**, and the System V Math Library, **libmsaa.a**. The **pow** subroutine in **libm.a** or **libmsaa.a** may be loaded in error unless the **libbsd.a** library is listed before the **libm.a** or **libmsaa.a** library on the command line.

The **gcd** subroutine returns the greatest common denominator of  $a$  and  $b$  in  $c$ , and the **invert** subroutine computes  $c$  such that  $a*c \bmod b=1$ , for  $a$  and  $b$  relatively prime.

The **ncmp** subroutine returns a negative, 0, or positive integer value when  $a$  is less than, equal to, or greater than  $b$ , respectively.

The **move** subroutine copies  $a$  to  $b$ . The **min** subroutine and **mout** subroutine do decimal input and output while the **omin** subroutine and **omout** subroutine do octal input and output. More generally, the **fmin** subroutine and **fmout** subroutine do decimal input and output using file  $f$ , and the **m\_in** subroutine and **m\_out** subroutine do inputs and outputs with arbitrary radix  $n$ . On input, records should have the form of strings of digits terminated by a new line; output records have a similar form.

Programs that use the multiple-precision arithmetic functions must link with the **libbsd.a** library.

Bases for input and output should be less than or equal to 10.

**pow** is also the name of a standard math library routine.

## Parameters

<i>Length</i>	Specifies the length of an integer.
<i>Value</i>	Specifies the initial value to be used in the routine.
<i>a</i>	Specifies the first operand of the multiple-precision routines.
<i>b</i>	Specifies the second operand of the multiple-precision routines.
<i>c</i>	Contains the integer result.
<i>f</i>	A pointer of the type <b>FILE</b> that points to input and output files used with input/output routines.
<i>m</i>	Indicates modulo.
<i>n</i>	Provides a value used to specify radix with <b>m_in</b> and <b>m_out</b> , power with <b>rpow</b> , and divisor with <b>sdiv</b> .
<i>q</i>	Contains the quotient obtained from <b>mdiv</b> .
<i>r</i>	Contains the remainder obtained from <b>mdiv</b> , <b>sdiv</b> , and <b>msqrt</b> .

## Error Codes

Error messages and core images are displayed as a result of illegal operations and running out of memory.

## Files

**/usr/lib/libbsd.a** Object code library.

## Related Information

The **bc** command, **dc** command.

## madvise Subroutine

### Purpose

Advises the system of expected paging behavior.

### Library

Standard C Library (**libc.a**).

### Syntax

```
#include <sys/types.h>
#include <sys/mman.h>
```

```
int madvise( addr, len, behav)
caddr_t addr;
size_t len;
int behav;
```

### Description

The **madvise** subroutine permits a process to advise the system about its expected future behavior in referencing a mapped file region or anonymous memory region.

The **madvise** subroutine has no functionality and is supported for compatibility only.

### Parameters

- addr* Specifies the starting address of the memory region. Must be a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.
- len* Specifies the length, in bytes, of the memory region. If the *len* value is not a multiple of page size as returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter, the length of the region will be rounded up to the next multiple of the page size.
- behav* Specifies the future behavior of the memory region. The following values for the *behav* parameter are defined in the **/usr/include/sys/mman.h** file:

Value	Paging Behavior Message
-------	-------------------------

**MADV\_NORMAL**

The system provides no further special treatment for the memory region.

**MADV\_RANDOM**

The system expects random page references to that memory region.

**MADV\_SEQUENTIAL**

The system expects sequential page references to that memory region.

**MADV\_WILLNEED**

The system expects the process will need these pages.

**MADV\_DONTNEED**

The system expects the process does not need these pages.

**MADV\_SPACEAVAIL**

The system will ensure that memory resources are reserved.

## Return Values

When successful, the **madvise** subroutine returns 0. Otherwise, it returns -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **madvise** subroutine is unsuccessful, the **errno** global variable can be set to one of the following values:

**EINVAL** The *behav* parameter is invalid.  
**ENOSPC** The *behav* parameter specifies **MADV\_SPACEAVAIL** and resources cannot be reserved.

## Related Information

The **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine, **sysconf** subroutine.

List of Memory Manipulation Services and Understanding Paging Space Programming Requirements in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## makecontext or swapcontext Subroutine

### Purpose

Modifies the context specified by *ucp*.

### Library

(**libc.a**)

### Syntax

```
#include <ucontext.h>
```

```
void makecontext (ucontext_t *ucp, (void *func) (), int argc, ...);  
int swapcontext (ucontext_t *oucp, const ucontext_t *ucp);
```

### Description

The **makecontext** subroutine modifies the context specified by *ucp*, which has been initialized using **getcontext** subroutine. When this context is resumed using **swapcontext** subroutine or **setcontext** subroutine, program execution continues by calling *func* parameter, passing it the arguments that follow *argc* in the **makecontext** subroutine.

Before a call is made to **makecontext** subroutine, the context being modified should have a stack allocated for it. The value of *argc* must match the number of integer argument passed to *func* parameter, otherwise the behavior is undefined.

The **uc\_link** member is used to determine the context that will be resumed when the context being modified by **makecontext** subroutine returns. The **uc\_link** member should be initialized prior to the call to **makecontext** subroutine.

The **swapcontext** subroutine function saves the current context in the context structure pointed to by *oucp* parameter and sets the context to the context structure pointed to by *ucp*.

### Parameters

*ucp* A pointer to a user structure.

*oucp* A pointer to a user structure.  
*func* A pointer to a function to be called when *ucp* is restored.  
*argc* The number of arguments being passed to *func* parameter.

## Return Values

On successful completion, **swapcontext** subroutine returns 0. Otherwise, a value of -1 is returned and **errno** is set to indicate the error.

-1 Not successful and the **errno** global variable is set to one of the following error codes.

## Error Codes

**ENOMEM** The *ucp* argument does not have enough stack left to complete the operation.

## Related Information

The **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutine, **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **wait** subroutine, **getcontext** (“getcontext or setcontext Subroutine” on page 357) subroutine, **sigaction** subroutine, and **sigprocmask** subroutine.

---

## matherr Subroutine

### Purpose

Math error handling function.

### Library

System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
int matherr (x)
struct exception *x;
```

### Description

The **matherr** subroutine is called by math library routines when errors are detected.

You can use **matherr** or define your own procedure for handling errors by creating a function named **matherr** in your program. Such a user-designed function must follow the same syntax as **matherr**. When an error occurs, a pointer to the exception structure will be passed to the user-supplied **matherr** function. This structure, which is defined in the **math.h** file, includes:

```
int type;
char *name;
double arg1, arg2, retval;
```

## Parameters

<i>type</i>	Specifies an integer describing the type of error that has occurred from the following list defined by the <b>math.h</b> file: <b>DOMAIN</b> Argument domain error <b>SING</b> Argument singularity <b>OVERFLOW</b> Overflow range error <b>UNDERFLOW</b> Underflow range error <b>TLOSS</b> Total loss of significance <b>PLOSS</b> Partial loss of significance.
<i>name</i>	Points to a string containing the name of the routine that caused the error.
<i>arg1</i>	Points to the first argument with which the routine was invoked.
<i>arg2</i>	Points to the second argument with which the routine was invoked.
<i>retval</i>	Specifies the default value that is returned by the routine unless the user's <b>matherr</b> function sets it to a different value.

## Return Values

If the user's **matherr** function returns a non-zero value, no error message is printed, and the **errno** global variable will not be set.

## Error Codes

If the function **matherr** is not supplied by the user, the default error-handling procedures, described with the math library routines involved, will be invoked upon error. In every case, the **errno** global variable is set to **EDOM** or **ERANGE** and the program continues.

## Related Information

The **bessel: j0, j1, jn, y0, y1, yn** (“bessel: j0, j1, jn, y0, y1, or yn Subroutine” on page 119) subroutine, **exp, expm1, log, log10, log1p, pow** (“exp, expf, or expl Subroutine” on page 245) subroutine, **lgamma** (“gamma Subroutine” on page 336) subroutine, **hypot, cabs** (“hypot, hypotf, or hypotl Subroutine” on page 528) subroutine, **sin, cos, tan, asin, acos, atan, atan2** subroutine, **sinh, cosh, tanh** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## MatchAllAuths, MatchAnyAuths, MatchAllAuthsList, or MatchAnyAuthsList Subroutine

### Purpose

Compare authorizations.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```

int MatchAllAuths(CommaListOfAuths)
char *CommaListOfAuths;
int MatchAllAuthsList(CommaListOfAuths, NSListOfAuths)
char *CommaListOfAuths;
char *NSListOfAuths;
int MatchAnyAuths(CommaListOfAuths)
char *CommaListOfAuths;
int MatchAnyAuthsList(CommaListOfAuths, NSListOfAuths)
char *CommaListOfAuths;
char *NSListOfAuths;

```

## Description

The **MatchAllAuthsList** subroutine compares the *CommaListOfAuths* against the *NSListOfAuths*. It returns a non-zero value if all the authorizations in *CommaListOfAuths* are contained in *NSListOfAuths*. The **MatchAllAuths** subroutine calls the **MatchAllAuthsList** subroutine passing in the results of the **GetUserAuths** subroutine in place of *NSListOfAuths*. If *NSListOfAuths* contains the OFF keyword, **MatchAllAuthsList** will return a zero value. If *NSListOfAuths* contains the ALL keyword and not the OFF keyword, **MatchAllAuthsList** will return a non-zero value.

The **MatchAnyAuthsList** subroutine compares the *CommaListOfAuths* against the *NSListOfAuths*. It returns a non-zero value if one or more of the authorizations in *CommaListOfAuths* are contained in *NSListOfAuths*. The **MatchAnyAuths** subroutine calls the **MatchAnyAuthsList** subroutine passing in the results of the **GetUserAuths** subroutine in place of *NSListOfAuths*. If *NSListOfAuths* contains the OFF keyword, **MatchAnyAuthsList** will return a zero value. If *NSListOfAuths* contains the ALL keyword and not the OFF keyword, **MatchAnyAuthsList** will return a non-zero value.

## Parameters

<i>CommaListOfAuths</i>	Specifies one or more authorizations, each separated by a comma.
<i>NSListOfAuths</i>	Specifies zero or more authorizations. Each authorization is null terminated. The last entry in the list must be a null string.

## Return Values

The subroutines return a non-zero value if a proper match was found. Otherwise, they will return zero. If an error occurs, the subroutines will return zero and set **errno** to indicate the error. If the subroutine returns zero and no error occurred, **errno** is set to zero.

---

## mblen Subroutine

### Purpose

Determines the length in bytes of a multibyte character.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```

int mblen( MbString, Number)
const char *MbString;
size_t Number;

```

## Description

The **mblen** subroutine determines the length, in bytes, of a multibyte character.

## Parameters

<i>Mbstring</i>	Points to a multibyte character string.
<i>Number</i>	Specifies the maximum number of bytes to consider.

## Return Values

The **mblen** subroutine returns 0 if the *MbString* parameter points to a null character. It returns -1 if a character cannot be formed from the number of bytes specified by the *Number* parameter. If *MbString* is a null pointer, 0 is returned.

## Related Information

The “mbslen Subroutine” on page 798, “mbstowcs Subroutine” on page 804, and “mbtowc Subroutine” on page 805.

Subroutines, Example Programs, and Libraries, in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mblen Subroutine

### Purpose

Get number of bytes in a character (restartable).

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <wchar.h>
size_t mblen (const char *s, size_t n, mbstate_t *ps )
```

### Description

If *s* is not a null pointer, **mblen** determines the number of bytes constituting the character pointed to by *s*. It is equivalent to:

```
mbstate_t internal;
mbrtowc(NULL, s, n, ps != NULL ? ps : &internal);
```

If *ps* is a null pointer, the **mblen** function uses its own internal **mbstate\_t** object, which is initialized at program startup to the initial conversion state. Otherwise, the **mbstate\_t** object pointed to by *ps* is used to completely describe the current conversion state of the associated character sequence. The implementation will behave as if no function defined in this specification calls **mblen**.

The behavior of this function is affected by the LC\_CTYPE category of the current locale.



## Return Values

The **mbrlen** function returns the first of the following that applies:

<b>0</b>	If the next <b>n</b> or fewer bytes complete the character that corresponds to the null wide-character
<b>positive</b>	If the next <b>n</b> or fewer bytes complete a valid character; the value returned is the number of bytes that complete the character.
<b>(size_t)-2</b>	If the next <b>n</b> bytes contribute to an incomplete but potentially valid character, and all <b>n</b> bytes have been processed. When <b>n</b> has at least the value of the <b>MB_CUR_MAX</b> macro, this case can only occur if <b>s</b> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).
<b>(size_t)-1</b>	If an encoding error occurs, in which case the next <b>n</b> or fewer bytes do not contribute to a complete and valid character. In this case, <b>EILSEQ</b> is stored in <b>errno</b> and the conversion state is undefined.

## Error Codes

The **mbrlen** function may fail if:

<b>EINVAL</b>	<b>ps</b> points to an object that contains an invalid conversion state.
<b>EILSEQ</b>	Invalid character sequence is detected.

## Related Information

The **mbsinit** (“mbsinit Subroutine” on page 797) subroutine, **mbrtowc** (“mbrtowc Subroutine”) subroutine.

---

## mbrtowc Subroutine

### Purpose

Convert a character to a wide-character code (restartable).

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <wchar.h>
size_t mbrtowc (wchar_t * pwc, const char * s, size_t n, mbstate_t * ps) ;
```

### Description

If **s** is a null pointer, the **mbrtowc** function is equivalent to the call:

```
mbrtowc(NULL, '', 1, ps)
```

In this case, the values of the arguments **pwc** and **n** are ignored.

If **s** is not a null pointer, the **mbrtowc** function inspects at most **n** bytes beginning at the byte pointed to by **s** to determine the number of bytes needed to complete the next character (including any shift sequences). If the function determines that the next character is completed, it determines the value of the corresponding wide-character and then, if **pwc** is not a null pointer, stores that value in the object pointed to by **pwc**. If the corresponding wide-character is the null wide-character, the resulting state described is the initial conversion state.

If **ps** is a null pointer, the **mbrtowc** function uses its own internal **mbstate\_t** object, which is initialized at program startup to the initial conversion state. Otherwise, the **mbstate\_t** object pointed to by **ps** is used to

completely describe the current conversion state of the associated character sequence. The implementation will behave as if no function defined in this specification calls **mbrtowc**.

The behavior of this function is affected by the `LC_CTYPE` category of the current locale.

## Return Values

The **mbrtowc** function returns the first of the following that applies:

<b>0</b>	If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character (which is the value stored).
<b>positive</b>	If the next <i>n</i> or fewer bytes complete a valid character (which is the value stored); the value returned is the number of bytes that complete the character.
<b>(size_t)-2</b>	If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed (no value is stored). When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).
<b>(size_t)-1</b>	If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character (no value is stored). In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.

## Error Codes

The **mbrtowc** function may fail if:

<b>EINVAL</b>	<i>ps</i> points to an object that contains an invalid conversion state.
<b>EILSEQ</b>	Invalid character sequence is detected.

## Related Information

The **mbsinit** (“mbsinit Subroutine” on page 797) subroutine.

---

## mbsadvance Subroutine

### Purpose

Advances to the next multibyte character.

**Note:** The **mbsadvance** subroutine is specific to the manufacturer. It is not defined in the POSIX, ANSI, or X/Open standards. Use of this subroutine may affect portability.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
```

```
char *mbsadvance ( S )  
const char *S;
```

### Description

The **mbsadvance** subroutine locates the next character in a multibyte character string. The `LC_CTYPE` category affects the behavior of the **mbsadvance** subroutine.

## Parameters

*S* Contains a multibyte character string.

## Return Values

If the *S* parameter is not a null pointer, the **mbsadvance** subroutine returns a pointer to the next multibyte character in the string pointed to by the *S* parameter. The character at the head of the string pointed to by the *S* parameter is skipped. If the *S* parameter is a null pointer or points to a null string, a null pointer is returned.

## Examples

To find the next character in a multibyte string, use the following:

```
#include <mbstr.h>
#include <locale.h>
#include <stdlib.h>
main()
{
    char *mbs, *pmb;
    (void) setlocale(LC_ALL, "");
    /*
    ** Let mbs point to the beginning of a multi-byte string.
    */
    pmb = mbs;
    while(pmb){
        pmb = mbsadvance(mbs);
        /* pmb points to the next multi-byte character
        ** in mbs */
    }
}
```

## Related Information

The **mbsinvalid** (“mbsinvalid Subroutine” on page 798) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbscat, mbscmp, or mbscpy Subroutine

### Purpose

Performs operations on multibyte character strings.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
char *mbscat(MbString1, MbString2)
char *MbString1, *MbString2;
int mbscmp(MbString1, MbString2)
char *MbString1, *MbString2;
```

```
char *mbscopy(MbString1, MbString2)
char *MbString1, *MbString2;
```

## Description

The **mbscat**, **mbscmp**, and **mbscopy** subroutines operate on null-terminated multibyte character strings.

The **mbscat** subroutine appends multibyte characters from the *MbString2* parameter to the end of the *MbString1* parameter, appends a null character to the result, and returns *MbString1*.

The **mbscmp** subroutine compares multibyte characters based on their collation weights as specified in the **LC\_COLLATE** category. The **mbscmp** subroutine compares the *MbString1* parameter to the *MbString2* parameter, and returns an integer greater than 0 if *MbString1* is greater than *MbString2*. It returns 0 if the strings are equivalent and returns an integer less than 0 if *MbString1* is less than *MbString2*.

The **mbscopy** subroutine copies multibyte characters from the *MbString2* parameter to the *MbString1* parameter and returns *MbString1*. The copy operation terminates with the copying of a null character.

## Related Information

The **mbnscat**, **mbnscmp**, **mbnscopy** (“mbnscat, mbsncmp, or mbsncpy Subroutine” on page 799) subroutine, **wcscat**, **wcscmp**, **wcscopy** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbchr Subroutine

### Purpose

Locates a character in a multibyte character string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
```

```
char *mbchr( MbString, MbCharacter)
char *MbString;
mbchar_t MbCharacter;
```

### Description

The **mbchr** subroutine locates the first occurrence of the value specified by the *MbCharacter* parameter in the string pointed to by the *MbString* parameter. The *MbCharacter* parameter specifies a multibyte character represented as an integer. The terminating null character is considered to be part of the string.

The **LC\_CTYPE** category affects the behavior of the **mbchr** subroutine.

## Parameters

<i>MbString</i>	Points to a multibyte character string.
<i>MbCharacter</i>	Specifies a multibyte character represented as an integer.

## Return Values

The **mbschr** subroutine returns a pointer to the value specified by the *MbCharacter* parameter within the multibyte character string, or a null pointer if that value does not occur in the string.

## Related Information

The “mbspbrk Subroutine” on page 800, “mbsrchr Subroutine” on page 801, “mbstomb Subroutine” on page 803, **wcschr** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbsinit Subroutine

### Purpose

Determine conversion object status.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <wchar.h>
int mbsinit (const mbstate_t * p) ;
```

### Description

If *ps* is not a null pointer, the **mbsinit** function determines whether the object pointed to by *ps* describes an initial conversion state.

The **mbstate\_t** object is used to describe the current conversion state from a particular character sequence to a wide-character sequence (or vice versa) under the rules of a particular setting of the LC\_CTYPE category of the current locale.

The initial conversion state corresponds, for a conversion in either direction, to the beginning of a new character sequence in the initial shift state. A zero valued **mbstate\_t** object is at least one way to describe an initial conversion state. A zero valued **mbstate\_t** object can be used to initiate conversion involving any character sequence, in any LC\_CTYPE category setting.

### Return Values

The **mbsinit** function returns non-zero if *ps* is a null pointer, or if the pointed-to object describes an initial conversion state; otherwise, it returns zero.

If an **mbstate\_t** object is altered by any of the functions described as restartable, and is then used with a different character sequence, or in the other conversion direction, or with a different LC\_CTYPE category setting than on earlier function calls, the behavior is undefined.

## Related Information

The “mbrlen Subroutine” on page 792, “mbrtowc Subroutine” on page 793, **wctomb** subroutine, “mbsrtowcs Subroutine” on page 802, **wcsrtombs** subroutine.

---

## mbsinvalid Subroutine

### Purpose

Validates characters of multibyte character strings.

**Note:** The **mbsinvalid** subroutine is specific to the manufacturer. It is not defined in the POSIX, ANSI, or X/Open standards. Use of this subroutine may affect portability.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
```

```
char *mbsinvalid ( S )  
const char *S;
```

### Description

The **mbsinvalid** subroutine examines the string pointed to by the *S* parameter to determine the validity of characters. The **LC\_CTYPE** category affects the behavior of the **mbsinvalid** subroutine.

### Parameters

*S*      Contains a multibyte character string.

### Return Values

The **mbsinvalid** subroutine returns a pointer to the byte following the last valid multibyte character in the *S* parameter. If all characters in the *S* parameter are valid multibyte characters, a null pointer is returned. If the *S* parameter is a null pointer, the behavior of the **mbsinvalid** subroutine is unspecified.

## Related Information

The “mbsadvance Subroutine” on page 794.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*

---

## mbslen Subroutine

### Purpose

Determines the number of characters (code points) in a multibyte character string.

**Note:** The **mbslen** subroutine is specific to the manufacturer. It is not defined in the POSIX, ANSI, or X/Open standards. Use of this subroutine may affect portability.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdlib.h>
```

```
size_t mbslen( MbString)  
char *mbs;
```

## Description

The **mbslen** subroutine determines the number of characters (code points) in a multibyte character string. The **LC\_CTYPE** category affects the behavior of the **mbslen** subroutine.

## Parameters

*MbString*        Points to a multibyte character string.

## Return Values

The **mbslen** subroutine returns the number of multibyte characters in a multibyte character string. It returns 0 if the *MbString* parameter points to a null character or if a character cannot be formed from the string pointed to by this parameter.

## Related Information

The **mblen** (“mblen Subroutine” on page 791) subroutine, **mbstowcs** (“mbstowcs Subroutine” on page 804) subroutine, **mbtowc** (“mbtowc Subroutine” on page 805) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbsncat, mbsncmp, or mbsncpy Subroutine

### Purpose

Performs operations on a specified number of null-terminated multibyte characters.

**Note:** These subroutines are specific to the manufacturer. They are not defined in the POSIX, ANSI, or X/Open standards. Use of these subroutines may affect portability.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <mbstr.h>
```

```
char *mbsncat(MbString1, MbString2, Number)  
char * MbString1, * MbString2;  
size_t Number;
```

```

int mbsncmp(MbString1, MbString2, Number)
char *MbString1, *MbString2;
size_t Number;

char *mbsncpy(MbString1, MbString2, Number)
char *MbString1, *MbString2;
size_t Number;

```

## Description

The **mbsncat**, **mbsncmp**, and **mbsncpy** subroutines operate on null-terminated multibyte character strings.

The **mbsncat** subroutine appends up to the specified maximum number of multibyte characters from the *MbString2* parameter to the end of the *MbString1* parameter, appends a null character to the result, and then returns the *MbString1* parameter.

The **mbsncmp** subroutine compares the collation weights of multibyte characters. The **LC\_COLLATE** category specifies the collation weights for all characters in a locale. The **mbsncmp** subroutine compares up to the specified maximum number of multibyte characters from the *MbString1* parameter to the *MbString2* parameter. It then returns an integer greater than 0 if *MbString1* is greater than *MbString2*. It returns 0 if the strings are equivalent. It returns an integer less than 0 if *MbString1* is less than *MbString2*.

The **mbsncpy** subroutine copies up to the value of the *Number* parameter of multibyte characters from the *MbString2* parameter to the *MbString1* parameter and returns *MbString1*. If *MbString2* is shorter than *Number* multi-byte characters, *MbString1* is padded out to *Number* characters with null characters.

## Parameters

<i>MbString1</i>	Contains a multibyte character string.
<i>MbString2</i>	Contains a multibyte character string.
<i>Number</i>	Specifies a maximum number of characters.

## Related Information

The “mbscat, mbscmp, or mbscpy Subroutine” on page 795, “mbscat, mbscmp, or mbscpy Subroutine” on page 795, “mbscat, mbscmp, or mbscpy Subroutine” on page 795, **wcsncat** subroutine, **wcsncmp** subroutine, **wcsncpy** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbspbrk Subroutine

### Purpose

Locates the first occurrence of multibyte characters or code points in a string.

**Note:** The **mbspbrk** subroutine is specific to the manufacturer. It is not defined in the POSIX, ANSI, or X/Open standards. Use of this subroutine may affect portability.

### Library

Standard C Library (**libc.a**)



## Syntax

```
#include <mbstr.h>
```

```
char *mbspbrk( MbString1, MbString2)  
char *MbString1, *MbString2;
```

## Description

The **mbspbrk** subroutine locates the first occurrence in the string pointed to by the *MbString1* parameter, of any character of the string pointed to by the *MbString2* parameter.

## Parameters

<i>MbString1</i>	Points to the string being searched.
<i>MbString2</i>	Pointer to a set of characters in a string.

## Return Values

The **mbspbrk** subroutine returns a pointer to the character. Otherwise, it returns a null character if no character from the string pointed to by the *MbString2* parameter occurs in the string pointed to by the *MbString1* parameter.

## Related Information

The “mbschr Subroutine” on page 796, “mbsrchr Subroutine,” “mbstomb Subroutine” on page 803, **wcspbrk** subroutine, **wcswcs** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbsrchr Subroutine

### Purpose

Locates a character or code point in a multibyte character string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
```

```
char *mbsrchr( MbString, MbCharacter)  
char *MbString;  
int MbCharacter;
```

### Description

The **mbschr** subroutine locates the last occurrence of the *MbCharacter* parameter in the string pointed to by the *MbString* parameter. The *MbCharacter* parameter is a multibyte character represented as an integer. The terminating null character is considered to be part of the string.

## Parameters

<i>MbString</i>	Points to a multibyte character string.
<i>MbCharacter</i>	Specifies a multibyte character represented as an integer.

## Return Values

The **mbsrchr** subroutine returns a pointer to the *MbCharacter* parameter within the multibyte character string. It returns a null pointer if *MbCharacter* does not occur in the string.

## Related Information

The “mbschr Subroutine” on page 796, “mbspbrk Subroutine” on page 800, “mbstomb Subroutine” on page 803, **wcsrchr** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*

---

## mbsrtowcs Subroutine

### Purpose

Convert a character string to a wide-character string (restartable).

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <wchar.h>
size_t mbsrtowcs ((wchar_t * dst, const char ** src, size_t len, mbstate_t * ps) ;
```

### Description

The **mbsrtowcs** function converts a sequence of characters, beginning in the conversion state described by the object pointed to by *ps*, from the array indirectly pointed to by **src** into a sequence of corresponding wide-characters. If *dst* is not a null pointer, the converted characters are stored into the array pointed to by *dst*. Conversion continues up to and including a terminating null character, which is also stored.

Conversion stops early in either of the following cases:

- When a sequence of bytes is encountered that does not form a valid character.
- When *len* codes have been stored into the array pointed to by **dst** (and *dst* is not a null pointer).

Each conversion takes place as if by a call to the **mbrtowc** function.

If *dst* is not a null pointer, the pointer object pointed to by *src* is assigned either a null pointer (if conversion stopped due to reaching a terminating null character) or the address just past the last character converted (if any). If conversion stopped due to reaching a terminating null character, and if *dst* is not a null pointer, the resulting state described is the initial conversion state.

If *ps* is a null pointer, the **mbsrtowcs** function uses its own internal **mbstate\_t** object, which is initialised at program startup to the initial conversion state. Otherwise, the **mbstate\_t** object pointed to by *ps* is used to completely describe the current conversion state of the associated character sequence. The implementation will behave as if no function defined in this specification calls **mbsrtowcs**.

The behavior of this function is affected by the LC\_CTYPE category of the current locale.

## Return Values

If the input conversion encounters a sequence of bytes that do not form a valid character, an encoding error occurs. In this case, the **mbsrtowcs** function stores the value of the macro EILSEQ in `errno` and returns **(size\_t)-1**; the conversion state is undefined. Otherwise, it returns the number of characters successfully converted, not including the terminating null (if any).

## Error Codes

The **mbsrtowcs** function may fail if:

**EINVAL**     `ps` points to an object that contains an invalid conversion state.  
**EILSEQ**     Invalid character sequence is detected.

## Related Information

The “mbsinit Subroutine” on page 797, “mbrtowc Subroutine” on page 793.

---

## mbstomb Subroutine

### Purpose

Extracts a multibyte character from a multibyte character string.

**Note:** The **mbstomb** subroutine is specific to the manufacturer. It is not defined in the POSIX, ANSI, or X/Open standards. Use of this subroutine may affect portability.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
```

```
mbchar_t mbstomb ( MbString )  
const char *MbString;
```

### Description

The **mbstomb** subroutine extracts the multibyte character pointed to by the *MbString* parameter from the multibyte character string. The **LC\_CTYPE** category affects the behavior of the **mbstomb** subroutine.

### Parameters

*MbString*     Contains a multibyte character string.

### Return Values

The **mbstomb** subroutine returns the code point of the multibyte character as a **mbchar\_t** data type. If an unusable multibyte character is encountered, a value of 0 is returned.

## Related Information

The “mbschr Subroutine” on page 796, “mbspbrk Subroutine” on page 800, “mbsrchr Subroutine” on page 801.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbstowcs Subroutine

### Purpose

Converts a multibyte character string to a wide character string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
size_t mbstowcs( WcString, String, Number)
wchar_t *WcString;
const char *String;
size_t Number;
```

### Description

The **mbstowcs** subroutine converts the sequence of multibyte characters pointed to by the *String* parameter to wide characters and places the results in the buffer pointed to by the *WcString* parameter. The multibyte characters are converted until a null character is reached or until the number of wide characters specified by the *Number* parameter have been processed.

### Parameters

<i>WcString</i>	Points to the area where the result of the conversion is stored.
<i>String</i>	Points to a multibyte character string.
<i>Number</i>	Specifies the maximum number of wide characters to be converted.

### Return Values

The **mbstowcs** subroutine returns the number of wide characters converted, not including a null terminator, if any. If an invalid multibyte character is encountered, a value of -1 is returned. The *WcString* parameter does not include a null terminator if the value *Number* is returned.

If *WcString* is a null wide character pointer, the **mbstowcs** subroutine returns the number of elements required to store the wide character codes in an array.

### Error Codes

The **mbstowcs** subroutine fails if the following occurs:

**EILSEQ** Invalid byte sequence is detected.

### Related Information

The “mblen Subroutine” on page 791, “mbslen Subroutine” on page 798, “mbtowc Subroutine” on page 805, **wcstombs** subroutine, **wctomb** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview for Programming and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbwidth Subroutine

### Purpose

Determines the number of multibyte character string display columns.

**Note:** The **mbwidth** subroutine is specific to this manufacturer. It is not defined in the POSIX, ANSI, or X/Open standards. Use of this subroutine may affect portability.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mbstr.h>
```

```
int mbwidth ( MbString, Number )  
const char *MbString;  
size_t Number;
```

### Description

The **mbwidth** subroutine determines the number of display columns required for a multibyte character string.

### Parameters

<i>MbString</i>	Contains a multibyte character string.
<i>Number</i>	Specifies the number of bytes to read from the <i>s</i> parameter.

### Return Values

The **mbwidth** subroutine returns the number of display columns that will be occupied by the *MbString* parameter if the number of bytes (specified by the *Number* parameter) read from the *MbString* parameter form valid multibyte characters. If the *MbString* parameter points to a null character, a value of 0 is returned. If the *MbString* parameter does not point to valid multibyte characters, -1 is returned. If the *MbString* parameter is a null pointer, the behavior of the **mbwidth** subroutine is unspecified.

### Related Information

The **wcwidth** subroutine, **wcwidth** subroutine.

National Language Support Overview in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## mbtowc Subroutine

### Purpose

Converts a multibyte character to a wide character.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdlib.h>
```

```
int mbtowc ( WideCharacter, String, Number)  
wchar_t *WideCharacter;  
const char *String;  
size_t Number;
```

## Description

The **mbtowc** subroutine converts a multibyte character to a wide character and returns the number of bytes of the multibyte character.

The **mbtowc** subroutine determines the number of bytes that comprise the multibyte character pointed to by the *String* parameter. It then converts the multibyte character to a corresponding wide character and, if the *WideCharacter* parameter is not a null pointer, places it in the location pointed to by the *WideCharacter* parameter. If the *WideCharacter* parameter is a null pointer, the **mbtowc** subroutine returns the number of converted bytes but does not change the *WideCharacter* parameter value. If the *WideCharacter* parameter returns a null value, the multibyte character is not converted.

## Parameters

<i>WideCharacter</i>	Specifies the location where a wide character is to be placed.
<i>String</i>	Specifies a multibyte character.
<i>Number</i>	Specifies the maximum number of bytes of a multibyte character.

## Return Values

The **mbtowc** subroutine returns a value of 0 if the *String* parameter is a null pointer. The subroutine returns a value of -1 if the bytes pointed to by the *String* parameter do not form a valid multibyte character before the number of bytes specified by the *Number* parameter (or fewer) have been processed. It then sets the **errno** global variable to indicate the error. Otherwise, the number of bytes comprising the multibyte character is returned.

## Error Codes

The **mbtowc** subroutine fails if the following occurs:

**EILSEQ** Invalid byte sequence is detected.

## Related Information

The “mblen Subroutine” on page 791, “mbslen Subroutine” on page 798, “mbstowcs Subroutine” on page 804, **wcstombs** subroutine, **wctomb** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## memccpy, memchr, memcmp, memcpy, memset or memmove Subroutine

### Purpose

Performs memory operations.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <memory.h>

void *memccpy (Target, Source, C, N)
void *Target;
const void *Source;
int C;
size_t N;

void *memchr ( S, C, N)
const void *S;
int C;
size_t N;

int memcmp (Target, Source, N)
const void *Target, *Source;
size_t N;

void *memcpy (Target, Source, N)
void *Target;
const void *Source;
size_t N;

void *memset (S, C, N)
void *S;
int C;
size_t N;

void *memmove (Target, Source, N)
void *Source;
const void *Target;
size_t N;
```

### Description

The **memory** subroutines operate on memory areas. A memory area is an array of characters bounded by a count. The **memory** subroutines do not check for the overflow of any receiving memory area. All of the **memory** subroutines are declared in the **memory.h** file.

The **memccpy** subroutine copies characters from the memory area specified by the *Source* parameter into the memory area specified by the *Target* parameter. The **memccpy** subroutine stops after the first character specified by the *C* parameter (converted to the **unsigned char** data type) is copied, or after *N* characters are copied, whichever comes first. If copying takes place between objects that overlap, the behavior is undefined.

The **memcmp** subroutine compares the first *N* characters as the **unsigned char** data type in the memory area specified by the *Target* parameter to the first *N* characters as the **unsigned char** data type in the memory area specified by the *Source* parameter.

The **memcpy** subroutine copies *N* characters from the memory area specified by the *Source* parameter to the area specified by the *Target* parameter and then returns the value of the *Target* parameter.

The **memset** subroutine sets the first *N* characters in the memory area specified by the *S* parameter to the value of character *C* and then returns the value of the *S* parameter.

Like the **memcpy** subroutine, the **memmove** subroutine copies *N* characters from the memory area specified by the *Source* parameter to the area specified by the *Target* parameter. However, if the areas of the *Source* and *Target* parameters overlap, the move is performed nondestructively, proceeding from right to left.

The **memccpy** subroutine is not in the ANSI C library.

## Parameters

<i>Target</i>	Points to the start of a memory area.
<i>Source</i>	Points to the start of a memory area.
<i>C</i>	Specifies a character to search.
<i>N</i>	Specifies the number of characters to search.
<i>S</i>	Points to the start of a memory area.

## Return Values

The **memccpy** subroutine returns a pointer to character *C* after it is copied into the area specified by the *Target* parameter, or a null pointer if the *C* character is not found in the first *N* characters of the area specified by the *Source* parameter.

The **memchr** subroutine returns a pointer to the first occurrence of the *C* character in the first *N* characters of the memory area specified by the *S* parameter, or a null pointer if the *C* character is not found.

The **memcmp** subroutine returns the following values:

<b>Less than 0</b>	If the value of the <i>Target</i> parameter is less than the values of the <i>Source</i> parameter.
<b>Equal to 0</b>	If the value of the <i>Target</i> parameter equals the value of the <i>Source</i> parameter.
<b>Greater than 0</b>	If the value of the <i>Target</i> parameter is greater than the value of the <i>Source</i> parameter.

## Related Information

The **swab** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## mincore Subroutine

### Purpose

Determines residency of memory pages.

### Library

Standard C Library (**libc.a**).



## Syntax

```
int mincore ( addr, len, * vec)
caddr_t addr;
size_t len;
char *vec;
```

## Description

The **mincore** subroutine returns the primary-memory residency status for regions created from calls made to the **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine. The status is returned as a character for each memory page in the range specified by the *addr* and *len* parameters. The least significant bit of each character returned is set to 1 if the referenced page is in primary memory. Otherwise, the bit is set to 0. The settings of the other bits in each character are undefined.

## Parameters

- addr* Specifies the starting address of the memory pages whose residency is to be determined. Must be a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.
- len* Specifies the length, in bytes, of the memory region whose residency is to be determined. If the *len* value is not a multiple of the page size as returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter, the length of the region is rounded up to the next multiple of the page size.
- vec* Specifies the character array where the residency status is returned. The system assumes that the character array specified by the *vec* parameter is large enough to encompass a returned character for each page specified.

## Return Values

When successful, the **mincore** subroutine returns 0. Otherwise, it returns -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **mincore** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

- EFAULT** A part of the buffer pointed to by the *vec* parameter is out of range or otherwise inaccessible.
- EINVAL** The *addr* parameter is not a multiple of the page size as returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.
- ENOMEM** Addresses in the (*addr*, *addr* + *len*) range are invalid for the address space of the process, or specify one or more pages that are not mapped.

## Related Information

The **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine, **sysconf** subroutine.

List of Memory Manipulation Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## MIO\_ao\_read64 Subroutine

### Purpose

Read asynchronously from a file through MIO library.

### Library

Modular I/O Library (**libmio.a**)

## Syntax

```
#include <libmio.h>

int MIO_aio_read64( FileDescriptor, aiocbp )
int FileDescriptor;
struct aiocb64 *aiocbp;
```

## Description

This subroutine is an entry point of the MIO library for the **Legacy AIO aio\_read64** subroutine. Use this subroutine to instrument your application with the MIO library. You can replace the **Legacy AIO aio\_read64 kernel I/O** subroutine with this equivalent MIO subroutine. See *Modular I/O in Performance management* for MIO library implementation.

Use this subroutine to read asynchronously from an open file specified by the *FileDescriptor* parameter. The *FileDescriptor* parameter results from an **MIO\_open64** subroutine.

## Parameters

The parameters are those of the corresponding standard POSIX system call `aio_read64`.

## Return Values

The return values are those of the corresponding standard POSIX system call `aio_read64`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `aio_read64`.

## Location

`/usr/lib/libmio.a`

## Related Information

The *Modular I/O in Performance management*.

The `aio_read64`, `MIO_open64`, `MIO_close`, `MIO_lseek64`, `MIO_write`, `MIO_ftruncate64`, `MIO_fstat64`, `MIO_fcntl`, `MIO_ffinfo`, and `MIO_fsync` subroutines.

---

## MIO\_aio\_suspend64 Subroutine

### Purpose

Suspend the calling process until one or more asynchronous I/O requests are completed.

### Library

Modular I/O Library (`libmio.a`)

### Syntax

```
#include <libmio.h>

int MIO_aio_suspend64( Count, aiocbplist )
int Count;
struct aiocb64 **aiocbplist;
```

## Description

This subroutine is an entry point of the MIO library for the **Legacy AIO aio\_suspend64** subroutine. Use this subroutine to instrument your application with the MIO library. You can replace the **Legacy AIO aio\_suspend64 kernel I/O** subroutine with this equivalent MIO subroutine. See *Modular I/O in Performance management* for the MIO library implementation.

The **aio\_suspend64** subroutine suspends the calling process until one or more of the *Count* parameter asynchronous I/O requests are completed or a signal interrupts the subroutine. Specifically, the **aio\_suspend64** subroutine handles requests associated with the aio control block (aiocb) structures pointed to by the *aiocbplist* parameter.

## Parameters

The parameters are those of the corresponding standard POSIX system call `aio_suspend64`.

## Return Values

The return values are those of the corresponding standard POSIX system call `aio_suspend64`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `aio_suspend64`.

## Location

`/usr/lib/libmio.a`

## Related Information

The Modular I/O in *Performance management*.

The **aio\_suspend64**, **MIO\_open64**, **MIO\_close**, **MIO\_lseek64**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstat64**, **MIO\_fcntl**, **MIO\_ffinfo**, and **MIO\_fsync** subroutines.

---

## MIO\_aio\_write64 Subroutine

### Purpose

Write asynchronously to a file through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>
```

```
int MIO_aio_write64( FileDescriptor, aiocbp )
int FileDescriptor; struct aiocb64 *aiocbp;
struct aiocb64 *aiocbp;
```

### Description

This subroutine is an entry point of the MIO library for the **Legacy AIO aio\_write64** subroutine. Use this subroutine to instrument your application with the MIO library. You can replace the **Legacy AIO aio\_write64 kernel I/O** subroutine with this equivalent MIO subroutine. See *Modular I/O in Performance management* for the MIO library implementation.

Use this subroutine to write asynchronously to an open file specified by the *FileDescriptor* parameter. The *FileDescriptor* parameter results from an **MIO\_open64** subroutine.

## Parameters

The parameters are those of the corresponding standard POSIX system call `aio_write64`.

## Return Values

The return values are those of the corresponding standard POSIX system call `aio_write64`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `aio_write64`.

## Location

`/usr/lib/libmio.a`

## Related Information

Modular I/O in *Performance management*.

The `aio_write64`, `MIO_open64`, `MIO_close`, `MIO_lseek64`, `MIO_write`, `MIO_ftruncate64`, `MIO_fstat64`, `MIO_fcntl`, `MIO_ffinfo`, and `MIO_fsync` subroutines.

---

## MIO\_close Subroutine

### Purpose

Close a file descriptor through the MIO library.

### Library

Modular I/O library (`libmio.a`)

### Syntax

```
#include <libmio.h>
```

```
int MIO_close (FileDescriptor)
```

```
int FileDescriptor;
```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **close kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to close a file with the *FileDescriptor* parameter through the Modular I/O (MIO) library. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

### Parameters

The parameters are those of the corresponding standard POSIX system call `close`.

### Return Values

The return values are those of the corresponding standard POSIX system call `close`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `close`.

## Standard Output

MIO library outputs are flushed on the `MIO_close` subroutine call in the `stats` file.

The following is the information found in the diagnostic output file. It contains debug information:

- If you set the `stats` option of the `trace` module (`trace/stats`), it runs diagnostics from the `trace` module.
- If you set the `stats` option of the `pf` module (`pf/stats`), it runs diagnostics from the `pf` module.
- If you set the `stats` option of the `recov` module (`recov/stats`), it runs diagnostics from the recovery trace.
- If you set the `nostats` option of the `trace` or the `pf` module, these diagnostics are suppressed.

The diagnostic file name is defined in the `MIO_STATS` environment variable if the `stats` option is set to the default value of `mioout`.

To separate the `trace`, `pf` or `recov` module diagnostics from other outputs, set the `stats` options to the following other file names:

- `trace/stats=<tracefile>`
- `pf/stats=<pf file>`
- `recov/stats=<recovfile>`

The `tracefile`, `pf file` and `recovfile` are templates for the file names of module diagnostics output. You can give file names for the output of the `trace`, `pf` or `recov` module diagnostics.

Standard output includes the following information:

**Header**, which contains the following information:

- Date
- Hostname
- Enabled or disabled AIO
- Program name
- MIO library version
- Environment variables

**Debug**, which contains the following information:

- The list of all the debug options
- The table of all of the modules' definitions if the `DEF` debug option is set
- Open request made to the `MIO_open64` subroutine if the `OPEN` debug option is set
- The modules invoked if the `MODULES` debug option is set

**Trace module diagnostic**, which contains the following information:

- Time if the `TIMESTAMP` debug option is set
- Trace on close or on intermediate interrupt
- Trace module position in `module_list`
- Processed file name
- Rate, which is the amount of data divided by the total time. The total time here means the cumulative amount of time spent beneath the `trace` module
- Demand rate, which is the amount of data divided by the length of time when the file is opened (including the time of opening and closing the file)

- The current (when tracing) file size and the maximum size of the file during this file processing
- File system information: file type and sector size
- Open mode and flags of the file
- For each subroutine, the following information is displayed:

name of the subroutine  
 count of calling of this subroutine  
 time of processing for this subroutine

- For read or write subroutines, the following information is displayed:

requested (requested size to read or write)  
 total (real size read or write: returned by AIX system call)  
 min (minimum size to read or write)  
 max (maximum size to read or write)

- For the seek subroutine, the following information is displayed:

the average seek delta (total seek delta/seek count)

- For the **aread** or **awrite** subroutine:

count, time and rate of transfer time including suspend, and read or write time

- For the **fcntl** subroutine, the number of pages is returned.

The following is an example of a **trace** diagnostic:

*date*

**Trace oncloseor intermediate:**

*previous module or calling program* <-> *next module:file name:*

*(total transferred bytes/total time)=rate*

**demand rate**=rate/s=*total transferred bytes/(close time-open time)*

**current size**=*actual size of the file*

**max\_size**=*max size of the file*

**mode**=*file open mode*

**FileSystemType**=*file system type given by fststat(stat\_b.f\_vfstype)*

**sector size**=*Minimum direct i/o transfer size*

**oflags**=*file open flags*

**open**    *open count*    *open time*

**fcntl**   *fcntl count*    *fcntl time*

**read**    *read count*    *read time*    *requested size*    *total size*    *minimum*    *maximum*

**aread**   *aread count*    *aread time*    *requested size*    *total size*    *minimum*    *maximum*

**suspend** *count time rate*

**write**   *write count*    *write time*    *requested size*    *total size*    *minimum*    *maximum*

**seek**    *seek count*    *seek time*    **average seek delta**

**size**

**page**    *fcntl page\_info count*

The following is a sample of a **trace** diagnostic:

MIO statistics file : Tue May 10 14:14:08 2005

hostname=host1 : with Legacy aio available

Program=example

MIO library libmio.a 3.0.0.60 AIX 5.1 32 bit addressing built

Apr 19 2005 15:08:17

MIO\_INSTALL\_PATH=

MIO\_STATS        =example.stats

MIO\_DEBUG        =OPEN

MIO\_FILES        = \*.dat [ trace/stats ]

MIO\_DEFAULTS    = trace/kbytes

MIO\_DEBUG OPEN =T

```

Opening file file.dat
modules[11]=trace/stats
=====
Trace close : program <-> aix : file.dat : (4800/0.04)=111538.02 kbytes/s
demand rate=42280.91 kbytes/s=4800/(0.12-0.01))
current size=0 max_size=1600
mode =0640 FileSystemType=JFS sector size=4096
oflags =0x302=RDWR CREAT TRUNC
open          1      0.00
write         100    0.02      1600      1600      16384      16384
read          200    0.02      3200      3200      16384      16384
seek          101    0.01 average seek delta=-48503
fcntl         1      0.00
trunc         1      0.01
close         1      0.00
size          100
=====

```

The following is a template of the **pf** module diagnostic:

```

pf close for<name of the file in the cache>
pf close for global or private cache <global cache number>
<nb_pg_compute>page of<page-size> <sector_size> bytes per sector
<nb_real_pg_not_pf>/<nb_pg_not_pf> pages not preread for write
<nb_unused_pf>unused prefetches out of<nb_start_pf>
prefetch=<nb_pg_to_pf>
<number> of write behind
<number> of page syncs forced by ill formed writes
<number> of pages retained over close
<unit> transferred / Number of requests
program --> <bytes written into the cache by parent>/
<number of write from parent>--> pf -->
<written out of the cache from the child>/<number of partial page written>
program --> <bytes read out of the cache by parent>/
<number of read from parent><-- pf <--
<bytes read in from child of the cache>/<number of page read from child>

```

The following is explanation of the terms in the **pf** module template:

- *nb\_pg\_compute*= number of page compute by cache\_size/ page size
- *nb\_real\_pg\_not\_pf*= real number page not prefetch because of *pfw* option (suppress number of page prefetch because sector not *valid*)
- *nb\_pg\_not\_pf*= page of unused prefetch
- *nb\_unused\_pf*= number of started prefetch
- *nb\_pg\_to\_pf*= number of page to prefetch

The following is a sample of the **pf** module diagnostic:

```

pf close for /home/user1/pthread/258/SM20182_0.SCR300
50 pages of 2097152 bytes 131072 bytes per sector
133/133 pages not preread for write
23 unused prefetches out of 242 : prefetch=2
95 write behinds
mbytes transferred / Number of requests
program --> 257/257 --> pf --> 257/131 --> aix
program <-- 269/269 <-- pf <-- 265/133 <-- aix

```

The following is the **recov** module output:

If open or write routine failed, the **recov** module, if set, is called. The **recov** module adds the following comments in the output file:

- The value of the *open\_command* option
- The value of the *command* option
- The **errno**
- The index of retry

The following is a sample of the **recov** module:

```
15:30:00
  recov : command=ls -l file=file.dat errno=28 try=0
  recov : failure : new_ret=-1
```

## Location

**/usr/lib/libmio.a**

## Related Information

The Modular I/O in *Performance management*.

The **close**, **MIO\_open64**, **MIO\_lseek64**, **MIO\_read**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstat64**, **MIO\_fcntl**, **MIO\_ffinfo**, **MIO\_fsync** subroutines.

---

## MIO\_fcntl Subroutine

### Purpose

Control open file descriptors through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>

int MIO_fcntl ( FileDescriptor, Command, Argument )

int FileDescriptor, Command, Argument;
```

### Description

This subroutine is an entry point of the MIO library, offering the same features as the **fcntl** subroutine. Use this subroutine to instrument your application with the MIO library. You can replace the **fcntl kernel I/O** subroutine with this equivalent MIO subroutine. See Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to perform controlling operations on the open file specified by the *FileDescriptor* parameter. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

### Parameters

The parameters are those of the corresponding standard POSIX system call **fcntl**.

### Return Values

The return values are those of the corresponding standard POSIX system call **fcntl**.



## Error Codes

The error codes are those of the corresponding standard POSIX system call `fcntl`.

## Location

`/usr/lib/libmio.a`

## Related Information

The Modular I/O in *Performance management*.

The `fcntl`, `MIO_open64`, `MIO_close`, `MIO_lseek64`, `MIO_write`, `MIO_ftruncate64`, `MIO_fstat64`, `MIO_ffinfo`, and `MIO_fsync` subroutines.

---

## MIO\_ffinfo Subroutine

### Purpose

Return file information through the MIO library.

### Library

Modular I/O library (`libmio.a`)

### Syntax

```
#include <libmio.h>

int MIO_ffinfo (FileDescriptor, Command, Buffer, Length)

int FileDescriptor;

int Command;

struct diocapbuf *Buffer;

int Length;
```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **ffinfo kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for MIO library implementation.

Use this subroutine to obtain specific file information for the open file referenced by the *FileDescriptor* parameter. The *FileDescriptor* parameter results from the `MIO_open64` subroutine.

### Parameters

The parameters are those of the corresponding standard POSIX system call `ffinfo`.

### Return Values

The return values are those of the corresponding standard POSIX system call `ffinfo`.

### Error Codes

The error codes are those of the corresponding standard POSIX system call `ffinfo`.

## Location

`/usr/lib/libmio.a`

## Related Information

The Modular I/O in *Performance management*.

The **ffinfo**, **MIO\_open64**, **MIO\_close**, **MIO\_lseek64**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstat64**, **MIO\_fcntl**, and **MIO\_fsync** subroutines.

---

## MIO\_fstat64 Subroutine

### Purpose

Provide information about a file through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>
```

```
int MIO_fstat64 (FileDescriptor, Buffer)
```

```
int FileDescriptor;
```

```
struct stat64 *Buffer;
```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **fstat64 kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to obtain information about the open file referenced by *FileDescriptor* parameter. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

### Parameters

The parameters are those of the corresponding standard POSIX system call `fstat64`.

### Return Values

The return values are those of the corresponding standard POSIX system call `fstat64`.

### Error Codes

The error codes are those of the corresponding standard POSIX system call `fstat64`.

### Location

/usr/lib/libmio.a

## Related Information

The Modular I/O in *Performance management*.

The **fstat64**, **MIO\_open64**, **MIO\_close**, **MIO\_lseek64**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fcntl**, **MIO\_ffinfo**, and **MIO\_fsync** subroutines.

---

## MIO\_fsync Subroutine

### Purpose

Save changes in a file to permanent storage through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>

int MIO_fsync (FileDescriptor)

int FileDescriptor;
```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **fsync kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to save to permanent storage all modified data in the specified range of the open file specified by the *FileDescriptor* parameter. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

### Parameters

The parameters are those of the corresponding standard POSIX system call fsync.

### Return Values

The return values are those of the corresponding standard POSIX system call fsync.

### Error Codes

The error codes are those of the corresponding standard POSIX system call fsync.

### Location

/usr/lib/libmio.a

### Related Information

The Modular I/O in *Performance management*.

The **fsync** , **MIO\_open64**, **MIO\_close**, **MIO\_lseek64**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstats64**, **MIO\_fcntl**, and **MIO\_ffinfo** subroutines.

---

## MIO\_ftruncate64 Subroutine

### Purpose

Change the length of regular files through the MIO library.

### Library

Modular I/O library (**libmio.a**)

## Syntax

```
#include <libmio.h>

int MIO_ftruncate64 (FileDescriptor, Length)

int FileDescriptor;

int64 Length;
```

## Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **ftruncate64 kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to change the length of the open file specified by the *FileDescriptor* parameter through Modular I/O (MIO) library. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

## Parameters

The parameters are those of the corresponding standard POSIX system call `ftruncate64`.

## Return Values

The return values are those of the corresponding standard POSIX system call `ftruncate64`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `ftruncate64`.

## Location

/usr/lib/libmio.a

## Related Information

The Modular I/O in *Performance management*.

The **ftruncate64**, **MIO\_open64**, **MIO\_close**, **MIO\_lseek64**, **MIO\_write**, **MIO\_fstat64**, **MIO\_fcntl**, **MIO\_ffinfo**, and **MIO\_fsync** subroutines.

---

## MIO\_lio\_listio64 Subroutine

### Purpose

Initiate a list of asynchronous I/O requests with a single call.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>

int MIO_lio_listio64 (Command, List, Nent, Eventp)
int Command;
struct liocb64 *List;
int Nent;
struct event *Eventp;
```

## Description

This subroutine is an entry point of the MIO library for the **Legacy AIO lio\_listio64** Subroutine. Use this subroutine to instrument your application with MIO library. You can replace the **Legacy AIO lio\_listio64 kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

The **lio\_listio64** subroutine allows the calling process to initiate the *Nent* parameter asynchronous I/O requests. These requests are specified in the *liocb* structures pointed to by the elements of the *List* array. The call may block or return immediately depending on the *Command* parameter. If the *Command* parameter requests that I/O completion be asynchronously notified, a SIGIO signal is delivered when all of the I/O operations are completed.

## Parameters

The parameters are those of the corresponding standard POSIX system call `lio_listio64`.

## Return Values

The return values are those of the corresponding standard POSIX system call `lio_listio64`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `lio_listio64`.

## Location

`/usr/lib/libmio.a`

## Related Information

The Modular I/O in *Performance management*.

The `lio_listio64`, `MIO_open64`, `MIO_close`, `MIO_lseek64`, `MIO_write`, `MIO_ftruncate64`, `MIO_fstat64`, `MIO_fcntl`, `MIO_ffinfo`, and `MIO_fsycn` subroutines.

---

## MIO\_lseek64 Subroutine

### Purpose

Move the read-write file pointer through the MIO library.

### Library

Modular I/O library (`libmio.a`)

### Syntax

```
#include <libmio.h>
```

```
int64 MIO_lseek64 (FileDescriptor, Offset, Whence)
int FileDescriptor;
int64 Offset;
int Whence;
```

## Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **fseek64 kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to set the read-write file pointer for the open file specified by the *FileDescriptor* parameter through the Modular I/O (MIO) library. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

## Parameters

The parameters are those of the corresponding standard POSIX system call **lseek64**.

## Return Values

The return values are those of the corresponding standard POSIX system call **lseek64**.

## Error Codes

The error codes are those of the corresponding standard POSIX system call **lseek64**.

## Location

/usr/lib/libmio.a

## Related Information

The Modular I/O in *Performance management*.

The **lseek64**, **MIO\_open64**, **MIO\_close**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstat64**, **MIO\_fcntl**, **MIO\_ffinfo**, and **MIO\_fsync** subroutines.

---

## MIO\_open64 Subroutine

### Purpose

Opens a file for reading or writing through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>
int MIO_open64 (Path, OFlag, Mode, Extra)
char *Path;
int OFlag;int Mode;
struct mio_extra *Extra;
```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **open64 kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to open a file through the Modular I/O (MIO) library. This library creates the context for this open file, according to the configuration set in MIO environment variables, or in the *Extra* parameter.

To analyze your application I/O and tune the I/O, use the MIO subroutines in the place of the standard I/O subroutines.

The MIO subroutines are:

- MIO\_close
- MIO\_lseek64
- MIO\_read
- MIO\_write
- MIO\_ftruncate64
- MIO\_fstat64
- MIO\_fcntl
- MIO\_ffinfo
- MIO\_fsync

The standard I/O subroutines are:

- close
- lseek64
- read
- write
- ftruncate64
- fstat64
- fcntl
- ffinfo
- fsync

## Parameters

*Extra*

Specifies some extra arguments for the MIO library. The simplest implementation is for any application to pass a zero pointer as the fourth argument. The fourth argument is a pointer to the `mio_extra` structure, you can usually pass a zero pointer, but you can also pass an `mio_extra` pointer (use this technique only if you are very familiar with how to code this argument).

The `mio_extra` structure is defined in the following way:

```
struct mio_extra {
    int cookie ;
    /* Default value: MIO_EXTRA_COOKIE/

    int taskid ;
    /* for later */

    int64 bufsiz ;
    /* if > 1 : force the prefetch for write pffw */

    char *modules ;
    /* explicit module name,
    if any modules returns from MIO_FILES environment variable match */

    char *logical_name ;
    /* logical file name to open
    if file name don't match with MIO_FILES regexp */

    int flags ;
    /* if MIO_EXTRA_SKIP_MIO_FILES_FLAG :
    don't use MIO_FILES env variable, but use extra->modules */
};
```

<i>Mode</i>	Specifies the modes. For more information, see the <i>Mode</i> flag in the <b>open64</b> subroutine.
<i>Oflag</i>	Specifies the type of access, the special open processing, the type of update, and the initial state of the open file. For more information, see the <b>open64</b> subroutine.
<i>Path</i>	Specifies the file to be opened.

**Note:** For applications that would not use the environment variable interface to apply the MIO modules to a file, the `mio_extra` hook provides an easy way to do that.

## Environment variables

MIO is controlled by the following environment variables, which define the MIO features and are processed by the **MIO\_open64** subroutine:

The *MIO\_STATS* variable is used to indicate a file that will be used as a repository for diagnostic messages and for output requested from the MIO modules. It is interpreted as a file name with two special cases. If the file is either the *stderr* or *stdout* output, the output will be directed towards the appropriate file stream. If the first character of the *MIO\_STATS* variable is a plus sign (+), the file name to be used is the string following the plus sign (+), and the file is opened for appending. Without the preceding plus sign (+), the file is overwritten.

The *MIO\_FILES* variable is the key to determine which modules are to be invoked for a given file when the **MIO\_open64** subroutine is called. The format for the *MIO\_FILES* variable is the following:

```
first_file_name_list [ module list ] second_file_name_list [ module list ] ...
```

When the **MIO\_open64** subroutine is called, MIO checks for the existence of the *MIO\_FILES* variable and parses it as follows:

The *MIO\_FILES* variable is parsed left to right. All characters up to the next occurrence of the bracket ([]) are taken as a file name list. A file name list is a colon-separated list of file name templates. A file name template is used to match the name of the file opened by MIO and can use the following wildcard characters:

- \* Matches zero or more characters of a directory or file name.
- ? Matches one character of a directory or file name.
- \*\* Matches all remaining characters of a full path name.

If the file name templates does not contain a forward slash (/), then all of the path directory information in the file name passed to the **MIO\_open64** subroutine is ignored and matching is applied only to the file name of the file being opened.

If the name of the file being opened is matched by one of the file name templates in the file name list then the module list to be invoked is taken as the string between brackets ([ ]). If the name of the file match two or more file name templates, the first match is taken into account. If the name of the file being opened does not match any of the file name templates in any of the file name lists then the file is opened with a default invocation of the AIX module.

If a match has occurred, the modules to be invoked are taken from the associated module list in the *MIO\_FILES* variable. The modules are invoked left to right, with the left-most being closest to the user program and the right-most being closest to the operating system. If the module list does not start with the MIO module, a default invocation of the MIO module is added as a prefix. If the AIX module is not specified, a default invocation of the AIX module is appended.

The following is an example of the *MIO\_FILES* variable:

```
setenv MIO_FILES " *.dat [ trace/stats ]"
```



Assume the *MIO\_FILES* variable is set as follows:

```
MIO_FILES= *.dat:*.scr [ trace ] *.f01:*.f02:*.f03 [ trace | pf | trace ]
```

If the **test.dat** file is opened by the **MIO\_open64** subroutine, the **test.dat** file name matches **\*.dat** and the following modules are invoked:

```
mio | trace | aix
```

If the **test.f02** file is opened by the **MIO\_open64** subroutine, the **test.f02** file name matches the second file name templates in the second file name list and the following modules are invoked:

```
mio | trace | pf | trace | aix
```

Each module has its own hardcoded default options for a default invocation. You can override the default options by specifying them in the associated *MIO\_FILES* module list. The following example turns on the **stats** option for the trace module and requests that the output be directed to the **my.stats** file:

```
MIO_FILES= *.dat : *.scr [ trace/stats=my.stats ]
```

The options for a module are delimited with a forward slash (/). Some options require an associated string value and others might require an integer value. For those requiring a string value, if the string includes a forward slash (/), enclose the string in braces ({ }).

For those options requiring an integer value, append the integer value with a k, m, g, or t to represent kilo, mega, giga, or tera. You might also input integer values in base 10, 8, or 16. If you add a 0x prefix to the integer value, the integer is interpreted as base 16. If you add a 0 prefix to the integer value, the integer is interpreted as base 8. If you add neither a 0x prefix nor a 0 prefix to the integer value, the integer is interpreted as base 10.

The *MIO\_DEFAULTS* variable is intended as a way to keep the *MIO\_FILES* variable more readable. If the user is specifying several modules for multiple file name list and module list pairs, then the *MIO\_FILES* variable might become quite long. To repeatedly override the hardcoded defaults in the same manner, you can specify new defaults for a module by specifying such defaults in the *MIO\_DEFAULTS* variable. The *MIO\_DEFAULTS* variable is a comma separated list of modules with their new defaults.

The following is an example of the *MIO\_DEFAULTS* variable:

```
setenv MIO_DEFAULTS " trace/kbytes "
```

Assume that *MIO\_DEFAULTS* variable is set as follows:

```
MIO_DEFAULTS = trace/events=prob.events , aix/debug
```

Any default invocation of the trace module will have binary event tracing enabled and directed towards the **prob.events** file and any default invocation of the AIX module will have debug enabled.

The *MIO\_DEBUG* variable is intended as an aid in debugging the use of MIO. MIO searches the *MIO\_DEFAULTS* variable for keywords and provides debugging output for the option. The available keywords are the following:

**ALL** Turns on all of the *MIO\_DEBUG* variable keywords.

**ENV** Outputs environment variable matching requests.

**OPEN** Outputs open requests made to the **MIO\_open64** subroutine.

#### **MODULES**

Outputs modules invoked for each call to the **MIO\_open64** subroutine.

#### **TIMESTAMP**

Places a timestamp preceding each entry into a **stats** file.

**DEF** Outputs the definition table of each module. When the file opens, the outputs of all of the MIO library's definitions are processed for all the MIO library modules.

## Return Values

The return values are those of the corresponding standard POSIX system call **open64**.

## Error Codes

The error codes are those of the corresponding standard POSIX system call **open64**.

## Standard Output

There is no MIO library output for the **MIO\_open64** subroutine.

**Note:** MIO library output statistics are written in the **MIO\_close** subroutine. This output filename is configurable with the *MIO\_STATS* environment variable.

In the **example.stats** MIO output file, the module trace is set and reported, and the open requests are output. All of the values are in kilobytes.

## Examples

The following **example.c** file issues 100 writes of 16 KB, seeks to the beginning of the file, issues 100 reads of 16 KB, and then seeks backward through the file reading 16 KB records. At the end the file is truncated to 0 bytes in length.

The *filename* argument to the following example is the file to be created, written to and read forwards and backwards:

```
-----
#define _LARGE_FILES
#include <fcntl.h>
#include <stdio.h>
#include <errno.h>

#include "libmio.h"

/* Define open64, lseek64 and ftruncate64, not
 * open, lseek, and ftruncate that are used in the code. This is
 * because libmio.h defines _LARGE_FILES which forces <fcntl.h> to
 * redefine open, lseek, and ftruncate as open64, lseek64, and
 * ftruncate64
 */

#define open64(a,b,c) MIO_open64(a,b,c,0)
#define close        MIO_close
#define lseek64      MIO_lseek64
#define write        MIO_write
#define read         MIO_read
#define ftruncate64  MIO_ftruncate64

#define RECSIZE 16384
#define NREC    100

main(int argc, char **argv)
{
    int i, fd, status ;
    char *name ;
    char *buffer ;
    int64 ret64 ;

    if( argc < 2 ){
        fprintf(stderr,"Usage : example file_name\n");
    }
}
```

```

    exit(-1);
}
name = argv[1] ;

buffer = (char *)malloc(RECSIZE);
memset( buffer, 0, RECSIZE ) ;

fd = open(name, O_RDWR|O_TRUNC|O_CREAT, 0640 ) ;
if( fd < 0 ){
    fprintf(stderr,"Unable to open file %s errno=%d\n",name,errno);
    exit(-1);
}

/* write the file */
for(i=0;i<NREC;i++){
    status = write( fd, buffer, RECSIZE ) ;
}

/* read the file forwards */
ret64 = lseek(fd, 0, SEEK_SET ) ;
for(i=0;i<NREC;i++){
    status = read( fd, buffer, RECSIZE ) ;
}

/* read the file backwards */
for(i=0;i<NREC;i++){
    ret64 = lseek(fd, (NREC-i-1)*RECSIZE, SEEK_SET ) ;
    status = read( fd, buffer, RECSIZE ) ;
}

/* truncate the file back to 0 bytes*/
status = ftruncate( fd, 0 ) ;

free(buffer);

/* close the file */
status = close(fd);
}

```

-----

Both a script that sets the environment variables, compiles and calls the application and the **example.c** example are delivered and installed with the **libmio** file, as follows:

```

cc -o example example.c -lmio

./example file.dat

```

The following environment variables are set to configure MIO:

```

setenv MIO_STATS example.stats
setenv MIO_FILES " *.dat [ trace/stats ] "
setenv MIO_DEFAULTS " trace/kbytes "
setenv MIO_DEBUG OPEN

```

See the **/usr/samples/libmio/README** file and sample files for details.

## Location

**/usr/lib/libmio.a**

## Related Information

The Modular I/O in *Performance management*.

The `open`, `MIO_close`, `MIO_lseek64`, `MIO_read`, `MIO_write`, `MIO_ftruncate64`, `MIO_fstat64`, `MIO_fcntl`, `MIO_ffinfo`, and `MIO_fsync` subroutines.

---

## MIO\_open Subroutine

### Purpose

Opens a file for reading or writing through the MIO library.

### Library

Modular I/O library (`libmio.a`)

### Syntax

```
#include <libmio.h>
int MIO_open (Path, OFlag, Mode, Extra)
char *Path;
int OFlag;
int Mode;
struct mio_extra *Extra;
```

### Description

The `MIO_open` subroutine is a redirection to the `MIO_open64` subroutine and is an entry point of the MIO library. To use the MIO library, the files have to be opened with the `O_LARGEFILE` flag. For more details on the `O_LARGEFILE` flag, see the `fcntl.h` File.

Use the `MIO_open` subroutine to instrument your application with the MIO library. You can replace the `open kernel I/O` subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to open a file through the Modular I/O (MIO) library. This library creates the context for this open file, according to the configuration set in the MIO environment variables, or in the `Extra` parameter.

To analyze your application I/O and tune the I/O, use the MIO subroutines in the place of the standard I/O subroutines.

The MIO subroutines are:

- `MIO_close`
- `MIO_lseek64`
- `MIO_read`
- `MIO_write`
- `MIO_ftruncate64`
- `MIO_fstat64`
- `MIO_fcntl`
- `MIO_ffinfo`
- `MIO_fsync`

The standard I/O subroutines are:

- `close`
- `lseek64`
- `read`
- `write`

- ftruncate64
- fstat64
- fcntl
- ffinfo
- fsync

## Parameters

*Extra* Specifies additional arguments for the MIO library. The simplest implementation is to pass a zero pointer as the fourth argument. The fourth argument is a pointer to the `mio_extra` structure, you can usually pass a zero pointer, but you can also pass an `mio_extra` pointer (use this technique only if you are very familiar with how to code this argument).

The `mio_extra` structure is defined as follows:

```
struct mio_extra {
    int cookie ;
    /* Default value: MIO_EXTRA_COOKIE/

    int taskid ;
    /* for later */

    int64 bufsiz ;
    /* if > 1 : force the prefetch for write pffw */

    char *modules ;
    /* explicit module name,
    if any modules returns from MIO_FILES environment variable match */

    char *logical_name ;
    /* logical file name to open
    if file name don't match with MIO_FILES regexp */

    int flags ;
    /* if MIO_EXTRA_SKIP_MIO_FILES_FLAG :
    don't use MIO_FILES env variable, but use extra->modules */
};
```

*Mode* Specifies the modes. For more information, see the *Mode* flag in the **open64** subroutine.

*Oflag* Specifies the type of access, the special open processing, the type of update, and the initial state of the open file. For more information, see the **open64** subroutine.

*Path* Specifies the file to be opened.

**Note:** For applications that would not use the environment variable interface to apply MIO modules to a file, the `mio_extra` hook provides an easy way to do that.

## Environment variables

MIO is controlled through the following four environment variables. These environment variables, which define the MIO features, are processed by the **MIO\_open64** subroutine.

The `MIO_STATS` variable is used to indicate a file that will be used as a repository for diagnostic messages and for output requested from the MIO modules. It is interpreted as a file name with two special cases. If the file is either `stderr` or `stdout` output, the output will be directed towards the appropriate file stream. If the first character of the `MIO_STATS` variable is a plus sign (+), the file name to be used is the string following the plus sign (+), and the file is opened for appending. Without the preceding plus sign (+), the file is overwritten.

The `MIO_FILES` variable is the key to determine which modules are to be invoked for a given file when the **MIO\_open64** subroutine is called. The format for the `MIO_FILES` variable is the following:

*first\_file\_name\_list* [ *module list* ] *second\_file\_name\_list* [ *module list* ]

When the **MIO\_open64** subroutine is called, MIO checks for the existence of the *MIO\_FILES* variable and parses it as follows:

The *MIO\_FILES* variable is parsed left to right. All characters up to the next occurrence of the bracket ([]) are taken as a file name list. A file name list is a colon-separated list of file name templates. A file name template is used to match the name of the file opened by MIO and can use the following wildcard characters:

- \* Matches zero or more characters of a directory or file name.
- ? Matches one character of a directory or file name.
- \*\* Matches all remaining characters of a full path name.

If the file name template does not contain a forward slash (/), then all of the path directory information in the file name passed to the **MIO\_open64** subroutine is ignored and matching is applied only to the file name of the file being opened.

If the name of the file being opened is matched by one of the file name templates in the file name list then the module list to be invoked is taken as the string between brackets ([ ]). If the name of the file match two or more file name templates, the first match is taken into account. If the name of the file being opened does not match any of the file name templates in any of the file name lists then the file is opened with a default invocation of the AIX module.

If a match has occurred, the modules to be invoked are taken from the associated module list in the *MIO\_FILES* variable. The modules are invoked left to right, with the left-most being closest to the user program and the right-most being closest to the operating system. If the module list does not start with the MIO module, a default invocation of the MIO module is added as a prefix. If the AIX module is not specified, a default invocation of the AIX module is appended.

The following is an example of the *MIO\_FILES* variable:

```
setenv MIO_FILES " *.dat [ trace/stats ]"
```

Assume the *MIO\_FILES* variable is set as follows:

```
MIO_FILES= *.dat:*.scr [ trace ] *.f01:*.f02:*.f03 [ trace | pf | trace ]
```

If the **test.dat** file is opened by the **MIO\_open64** subroutine, the **test.dat** file name matches **\*.dat** and the following modules are invoked:

```
mio | trace | aix
```

If the **test.f02** file is opened by the **MIO\_open64** subroutine, the **test.f02** file name matches the second file name templates in the second file name list and the following modules are invoked:

```
mio | trace | pf | trace | aix
```

Each module has its own hardcoded default options for a default invocation. You can override the default options by specifying them in the associated *MIO\_FILES* module list. The following example turns on the **stats** option for the trace module and requests that the output be directed to the **my.stats** file:

```
MIO_FILES= *.dat : *.scr [ trace/stats=my.stats ]
```

The options for a module are delimited with a forward slash (/). Some options require an associated string value and others might require an integer value. For those requiring a string value, if the string includes a forward slash (/), enclose the string in braces ({}).

For those options requiring an integer value, append the integer value with a k, m, g, or t to represent kilo, mega, giga, or tera. You might also input integer values in base 10, 8, or 16. If you add a 0x prefix to the

integer value, the integer is interpreted as base 16. If you add a 0 prefix to the integer value, the integer is interpreted as base 8. If you add neither a 0x prefix nor a 0 prefix to the integer value, the integer is interpreted as base 10.

The *MIO\_DEFAULTS* variable is intended as a way to keep the *MIO\_FILES* variable more readable. If the user is specifying several modules for multiple file name list and module list pairs, then the *MIO\_FILES* variable might become quite long. To repeatedly override the hardcoded defaults in the same manner, you can specify new defaults for a module by specifying such defaults in the *MIO\_DEFAULTS* variable. The *MIO\_DEFAULTS* variable is a comma separated list of modules with their new defaults.

The following is an example of the *MIO\_DEFAULTS* variable:

```
setenv MIO_DEFAULTS " trace/kbytes "
```

Assume that *MIO\_DEFAULTS* variable is set as follows:

```
MIO_DEFAULTS = trace/events=prob.events , aix/debug
```

Any default invocation of the trace module will have binary event tracing enabled and directed towards the **prob.events** file and any default invocation of the AIX module will have debug enabled.

The *MIO\_DEBUG* variable is intended as an aid in debugging the use of MIO. MIO searches the *MIO\_DEFAULTS* variable for keywords and provides debugging output for the option. The available keywords are the following:

**ALL** Turns on all of the *MIO\_DEBUG* variable keywords.

**ENV** Outputs environment variable matching requests.

**OPEN** Outputs open requests made to the **MIO\_open64** subroutine.

#### **MODULES**

Outputs modules invoked for each call to the **MIO\_open64** subroutine.

#### **TIMESTAMP**

Places a timestamp preceding each entry into a **stats** file.

**DEF** Outputs the definition table of each module. When the file opens, the outputs of all of the MIO library's definitions are processed for all the MIO library modules.

## **Return values**

The return values are those of the corresponding standard POSIX system call **open64**.

## **Error codes**

The error codes are those of the corresponding standard POSIX system call **open64**.

## **Standard output**

There is no MIO library output for the **MIO\_open64** subroutine.

**Note:** MIO library output statistics are written in the **MIO\_close** subroutine. This output filename is configurable with the *MIO\_STATS* environment variable.

In the **example.stats**. MIO output file, the module trace is set and reported, and the open requests are output. All the values are in kilobytes.

## **Examples**

The following **example.c** file issues 100 writes of 16 KB, seeks to the beginning of the file, issues 100 reads of 16 KB, and then seeks backward through the file reading 16 KB records. At the end the file is truncated to 0 bytes in length.

The *filename* argument to the following example is the file to be created, written to and read forwards and backwards:

```
-----
#define _LARGE_FILES
#include <fcntl.h>
#include <stdio.h>
#include <errno.h>

#include "libmio.h"

/* Define open64, lseek64 and ftruncate64, not
 * open, lseek, and ftruncate that are used in the code. This is
 * because libmio.h defines _LARGE_FILES which forces <fcntl.h> to
 * redefine open, lseek, and ftruncate as open64, lseek64, and
 * ftruncate64
 */

#define open64(a,b,c) MIO_open64(a,b,c,0)
#define close        MIO_close
#define lseek64      MIO_lseek64
#define write        MIO_write
#define read         MIO_read
#define ftruncate64  MIO_ftruncate64

#define RECSIZE 16384
#define NREC    100

main(int argc, char **argv)
{
    int i, fd, status ;
    char *name ;
    char *buffer ;
    int64 ret64 ;

    if( argc < 2 ){
        fprintf(stderr,"Usage : example file_name\n");
        exit(-1);
    }
    name = argv[1] ;

    buffer = (char *)malloc(RECSIZE);
    memset( buffer, 0, RECSIZE ) ;

    fd = open(name, O_RDWR|O_TRUNC|O_CREAT, 0640 ) ;
    if( fd < 0 ){
        fprintf(stderr,"Unable to open file %s errno=%d\n",name,errno);
        exit(-1);
    }

    /* write the file */
    for(i=0;i<NREC;i++){
        status = write( fd, buffer, RECSIZE ) ;
    }

    /* read the file forwards */
    ret64 = lseek(fd, 0, SEEK_SET ) ;
    for(i=0;i<NREC;i++){
        status = read( fd, buffer, RECSIZE ) ;
    }

    /* read the file backwards */
    for(i=0;i<NREC;i++){
        ret64 = lseek(fd, (NREC-i-1)*RECSIZE, SEEK_SET ) ;
        status = read( fd, buffer, RECSIZE ) ;
    }

    /* truncate the file back to 0 bytes*/

```



```

    status = ftruncate( fd, 0 ) ;

    free(buffer);

/* close the file */
    status = close(fd);
}

```

Both a script that sets the environment variables, compiles and calls the application and the **example.c** example are delivered and installed with the **libmio**, as follows:

```

cc -o example example.c -lmio

./example file.dat

```

The following environment variables are set to configure MIO:

```

setenv MIO_STATS example.stats
setenv MIO_FILES " *.dat [ trace/stats ] "
setenv MIO_DEFAULTS " trace/kbytes "
setenv MIO_DEBUG OPEN

```

See the **/usr/samples/libmio/README** and sample files for details.

## Location

**/usr/lib/libmio.a**

## Related Information

The Modular I/O in *Performance management*.

The **open**, **MIO\_close**, **MIO\_lseek64**, **MIO\_read**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstat64**, **MIO\_fcntl**, **MIO\_ffinfo**, **MIO\_fsync** subroutines.

---

## MIO\_read Subroutine

### Purpose

Read from a file through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```

#include <libmio.h>

int MIO_read(FileDescriptor,
Buffer, NBytes)
int FileDescriptor;
void * Buffer;
int NBytes;

```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **read kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to read to the number of bytes of data specified by the *NBytes* parameter from the file associated with the *FileDescriptor* parameter into the buffer, through the Modular I/O (MIO) library. The *Buffer* parameter points to the buffer. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

## Parameters

The parameters are those of the corresponding standard POSIX system call `read`.

## Return Values

The return values are those of the corresponding standard POSIX system call `read`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `read`.

## Location

`/usr/lib/libmio.a`

## Related Information

The Modular I/O in *Performance management*.

The `read`, **MIO\_open64**, **MIO\_close**, **MIO\_lseek64**, **MIO\_write**, **MIO\_ftruncate64**, **MIO\_fstat64**, **MIO\_fcntl**, **MIO\_ffinfo**, and **MIO\_fsync** subroutines.

---

## MIO\_write Subroutine

### Purpose

Write to a file through the MIO library.

### Library

Modular I/O library (**libmio.a**)

### Syntax

```
#include <libmio.h>

int MIO_write(FileDescriptor,
Buffer, NBytes)
int FileDescriptor;
void * Buffer;
int NBytes;
```

### Description

This subroutine is an entry point of the MIO library. Use this subroutine to instrument your application with the MIO library. You can replace the **write kernel I/O** subroutine with this equivalent MIO subroutine. See the Modular I/O in *Performance management* for the MIO library implementation.

Use this subroutine to write the number of bytes of data specified by the *NBytes* parameter from the buffer to the file associated with the *FileDescriptor* parameter through the Modular I/O (MIO) library. The *Buffer* parameter points to the buffer. The *FileDescriptor* parameter results from the **MIO\_open64** subroutine.

### Parameters

The parameters are those of the corresponding standard POSIX system call `write`.

## Return Values

The return values are those of the corresponding standard POSIX system call `write`.

## Error Codes

The error codes are those of the corresponding standard POSIX system call `write`.

## Location

`/usr/lib/libmio.a`

## Related Information

The Modular I/O in *Performance management*.

The `write`, `MIO_open64`, `MIO_close`, `MIO_lseek64`, `MIO_ftruncate64`, `MIO_fstat64`, `MIO_fcntl`, `MIO_ffinfo`, and `MIO_fsync` subroutines.

---

## mkdir Subroutine

### Purpose

Creates a directory.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/stat.h>
```

```
int mkdir ( Path, Mode)
const char *Path;
mode_t Mode;
```

### Description

The `mkdir` subroutine creates a new directory.

The new directory has the following:

- The owner ID is set to the process-effective user ID.
- If the parent directory has the *SetFileGroupID* (**S\_ISGID**) attribute set, the new directory inherits the group ID of the parent directory. Otherwise, the group ID of the new directory is set to the effective group ID of the calling process.
- Permission and attribute bits are set according to the value of the *Mode* parameter, with the following modifications:
  - All bits set in the process-file mode-creation mask are cleared.
  - The *SetFileUserID* and *Sticky* (**S\_ISVTX**) attributes are cleared.
- If the *Path* variable names a symbolic link, the link is followed. The new directory is created where the variable pointed.

## Parameters

<i>Path</i>	Specifies the name of the new directory. If Network File System (NFS) is installed on your system, this path can cross into another node. In this case, the new directory is created at that node.
<i>Mode</i>	To execute the <b>mkdir</b> subroutine, a process must have search permission to get to the parent directory of the <i>Path</i> parameter as well as write permission in the parent directory itself. Specifies the mask for the read, write, and execute flags for owner, group, and others. The <i>Mode</i> parameter specifies directory permissions and attributes. This parameter is constructed by logically ORing values described in the <b>sys/mode.h</b> file.

## Return Values

Upon successful completion, the **mkdir** subroutine returns a value of 0. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **mkdir** subroutine is unsuccessful and the directory is not created if one or more of the following are true:

<b>EACCES</b>	Creating the requested directory requires writing in a directory with a mode that denies write permission.
<b>EEXIST</b>	The named file already exists.
<b>EROFS</b>	The named file resides on a read-only file system.
<b>ENOSPC</b>	The file system does not contain enough space to hold the contents of the new directory or to extend the parent directory of the new directory.
<b>EMLINK</b>	The link count of the parent directory exceeds the maximum ( <b>LINK_MAX</b> ) number. ( <b>LINK_MAX</b> ) is defined in <b>limits.h</b> file.
<b>ENAMETOOLONG</b>	The <i>Path</i> parameter or a path component is too long and cannot be truncated.
<b>ENOENT</b>	A component of the path prefix does not exist or the <i>Path</i> parameter points to an empty string.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>EDQUOT</b>	The directory in which the entry for the new directory is being placed cannot be extended, or an i-node or disk blocks could not be allocated for the new directory because the user's or group's quota of disk blocks or i-nodes on the file system containing the directory is exhausted.

The **mkdir** subroutine can be unsuccessful for other reasons. See "Appendix A. Base Operating System Error Codes for Services That Require Path-Name Resolution" for a list of additional errors.

If NFS is installed on the system, the **mkdir** subroutine is also unsuccessful if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **chmod** ("chmod or fchmod Subroutine" on page 148) subroutine, **mknod** ("mknod or mkfifo Subroutine" on page 837) subroutine, **rmdir** subroutine, **umask** subroutine.

The **chmod** command, **mkdir** command, **mknod** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## mknod or mkfifo Subroutine

### Purpose

Creates an ordinary file, first-in-first-out (FIFO), or special file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/stat.h>
```

```
int mknod (const char * Path, mode_t Mode, dev_t Device)
char *Path;
int Mode;
dev_t Device;
int mkfifo (const char *Path, mode_t Mode)
const char *Path;
int Mode;
```

### Description

The **mknod** subroutine creates a new regular file, special file, or FIFO file. Using the **mknod** subroutine to create file types (other than FIFO or special files) requires root user authority.

For the **mknod** subroutine to complete successfully, a process must have both search and write permission in the parent directory of the *Path* parameter.

The **mkfifo** subroutine is an interface to the **mknod** subroutine, where the new file to be created is a FIFO or special file. No special system privileges are required.

The new file has the following characteristics:

- File type is specified by the *Mode* parameter.
- Owner ID is set to the effective user ID of the process.
- Group ID of the file is set to the group ID of the parent directory if the *SetGroupID* attribute (**S\_ISGID**) of the parent directory is set. Otherwise, the group ID of the file is set to the effective group ID of the calling process.
- Permission and attribute bits are set according to the value of the *Mode* parameter. All bits set in the file-mode creation mask of the process are cleared.

Upon successful completion, the **mkfifo** subroutine marks for update the *st\_atime*, *st\_ctime*, and *st\_mtime* fields of the file. It also marks for update the *st\_ctime* and *st\_mtime* fields of the directory that contains the new entry.

If the new file is a character special file having the **S\_IMPX** attribute (multiplexed character special file), when the file is used, additional path-name components can appear after the path name as if it were a directory. The additional part of the path name is available to the device driver of the file for interpretation. This feature provides a multiplexed interface to the device driver.

### Parameters

*Path* Names the new file. If Network File System (NFS) is installed on your system, this path can cross into another node.

<i>Mode</i>	Specifies the file type, attributes, and access permissions. This parameter is constructed by logically ORing values described in the <b>sys/mode.h</b> file.
<i>Device</i>	Specifies the ID of the device, which corresponds to the <code>st_rdev</code> member of the structure returned by the <b>statx</b> subroutine. This parameter is configuration-dependent and used only if the <i>Mode</i> parameter specifies a block or character special file. If the file you specify is a remote file, the value of the <i>Device</i> parameter must be meaningful on the node where the file resides.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **mknod** subroutine fails and the new file is not created if one or more of the following are true:

<b>EEXIST</b>	The named file exists.
<b>EDQUOT</b>	The directory in which the entry for the new file is being placed cannot be extended, or an i-node could not be allocated for the file because the user's or group's quota of disk blocks or i-nodes on the file system is exhausted.
<b>EISDIR</b>	The <i>Mode</i> parameter specifies a directory. Use the <b>mkdir</b> subroutine instead.
<b>ENOSPC</b>	The directory that would contain the new file cannot be extended, or the file system is out of file-allocation resources.
<b>EPERM</b>	The <i>Mode</i> parameter specifies a file type other than <b>S_IFIFO</b> , and the calling process does not have root user authority.
<b>EROFS</b>	The directory in which the file is to be created is located on a read-only file system.

The **mknod** and **mkfifo** subroutine can be unsuccessful for other reasons. See "Appendix. A Base Operating System Error Codes for Services That Require Path-Name Resolution" (Appendix A, "Base Operating System Error Codes for Services That Require Path-Name Resolution," on page 1345) for a list of additional errors.

If NFS is installed on the system, the **mknod** subroutine can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **chmod** ("chmod or fchmod Subroutine" on page 148) subroutine, **mkdir** ("mkdir Subroutine" on page 835) subroutine, **open**, **openx**, or **creat** ("open, openx, open64, creat, or creat64 Subroutine" on page 935) subroutine, **statx** subroutine, **umask** subroutine.

The **chmod** command, **mkdir** command, **mknod** command.

The **mode.h** file, **types.h** file.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## mktemp or mkstemp Subroutine

### Purpose

Constructs a unique file name.

## Libraries

Standard C Library (**libc.a**)

Berkeley Compatibility Library (**libbsd.a**)

## Syntax

```
#include <stdlib.h>
```

```
char *mktemp ( Template )  
char *Template;
```

```
int mkstemp ( Template )  
char *Template;
```

## Description

The **mktemp** subroutine replaces the contents of the string pointed to by the *Template* parameter with a unique file name.

**Note:** The **mktemp** subroutine creates a filename and checks to see if the file exist. If that file does not exist, the name is returned. If the user calls **mktemp** twice without creating a file using the name returned by the first call to **mktemp**, then the second **mktemp** call may return the same name as the first **mktemp** call since the name does not exist.

To avoid this, either create the file after calling **mktemp** or use the **mkstemp** subroutine. The **mkstemp** subroutine creates the file for you.

To get the BSD version of this subroutine, compile with Berkeley Compatibility Library (**libbsd.a**).

The **mkstemp** subroutine performs the same substitution to the template name and also opens the file for reading and writing.

In BSD systems, the **mkstemp** subroutine was intended to avoid a race condition between generating a temporary name and creating the file. Because the name generation in the operating system is more random, this race condition is less likely. BSD returns a file name of / (slash).

Former implementations created a unique name by replacing X's with the process ID and a unique letter.

## Parameters

*Template* Points to a string to be replaced with a unique file name. The string in the *Template* parameter is a file name with up to six trailing X's. Since the system randomly generates a six-character string to replace the X's, it is recommended that six trailing X's be used.

## Return Values

Upon successful completion, the **mktemp** subroutine returns the address of the string pointed to by the *Template* parameter.

If the string pointed to by the *Template* parameter contains no X's, and if it is an existing file name, the *Template* parameter is set to a null character, and a null pointer is returned; if the string does not match any existing file name, the exact string is returned.

Upon successful completion, the **mkstemp** subroutine returns an open file descriptor. If the **mkstemp** subroutine fails, it returns a value of -1.

## Related Information

The **getpid** (“getpid, getpgrp, or getppid Subroutine” on page 407) subroutine, **tmpfile** subroutine, **tmpnam** or **tempnam** subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## mlock and munlock Subroutine

### Purpose

Locks or unlocks a range of process address space.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
int mlock (addr, len)
const void *addr;
size_t len;
```

```
int munlock (addr, len)
const void *addr;
size_t len;
```

### Description

The **mlock** subroutine causes those whole pages containing any part of the address space of the process starting at address *addr* and continuing for *len* bytes to be memory-resident until unlocked or until the process exits or executes another process image. If the starting address *addr* is not a multiple of **PAGESIZE**, it is rounded down to the lowest page boundary. The *len* is rounded up to a multiple of **PAGESIZE**.

The **munlock** subroutine unlocks those whole pages containing any part of the address space of the process starting at address *addr* and continuing for *len* bytes, regardless of how many times **mlock** has been called by the process for any of the pages in the specified range.

If any of the pages in the range specified in a call to the **munlock** subroutine are also mapped into the address spaces of other processes, any locks established on those pages by another process are unaffected by the call of this process to the **munlock** subroutine. If any of the pages in the range specified by a call to the **munlock** subroutine are also mapped into other portions of the address space of the calling process outside the range specified, any locks established on those pages through other mappings are also unaffected by this call.

Upon successful return from **mlock**, pages in the specified range are locked and memory-resident. Upon successful return from **munlock**, pages in the specified range are unlocked with respect to the address space of the process.

The calling process must have the root user authority to use this subroutine.

### Parameters

*addr*                Specifies the address space of the process to be locked or unlocked.  
*len*                 Specifies the length in bytes of the address space.



## Return Values

Upon successful completion, the **mlock** and **munlock** subroutines return zero. Otherwise, no change is made to any locks in the address space of the process, the subroutines return -1 and set **errno** to indicate the error.

## Error Codes

The **mlock** and **munlock** subroutines fail if:

<b>ENOMEM</b>	Some or all of the address range specified by the <i>addr</i> and <i>len</i> parameters does not correspond to valid mapped pages in the address space of the process.
<b>EINVAL</b>	The process has already some plocked memory or the <i>len</i> parameter is negative.
<b>EPERM</b>	The calling process does not have the appropriate privilege to perform the requested operation.

The **mlock** subroutine might fail if:

<b>ENOMEM</b>	Locking the pages mapped by the specified range would exceed the limit on the amount of memory the process may lock.
---------------	--

## Related Information

“exec: execl, execl, execlp, execv, execve, execvp, or exec Subroutine” on page 236, “exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243, “fork, f\_fork, or vfork Subroutine” on page 291, “mlockall and munlockall Subroutine,” and “munmap Subroutine” on page 893.

---

## mlockall and munlockall Subroutine

### Purpose

Locks or unlocks the address space of a process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
int mlockall (flags)
int flags;
```

```
int munlockall (void);
```

### Description

The **mlockall** subroutine causes all of the pages mapped by the address space of a process to be memory-resident until unlocked or until the process exits or executes another process image. The *flags* parameter determines whether the pages to be locked are those currently mapped by the address space of the process, those that are mapped in the future, or both. The *flags* parameter is constructed from the bitwise-inclusive OR of one or more of the following symbolic constants, defined in the **sys/mman.h** header file:

#### **MCL\_CURRENT**

Lock all of the pages currently mapped into the address space of the process.

## MCL\_FUTURE

Lock all of the pages that become mapped into the address space of the process in the future, when those mappings are established.

When **MCL\_FUTURE** is specified, the future mapping functions might fail if the system is not able to lock this amount of memory because of lack of resources, for example.

The **munlockall** subroutine unlocks all currently mapped pages of the address space of the process. Any pages that become mapped into the address space of the process after a call to the **munlockall** subroutine are not locked, unless there is an intervening call to the **mlockall** subroutine specifying **MCL\_FUTURE** or a subsequent call to the **mlockall** subroutine specifying **MCL\_CURRENT**. If pages mapped into the address space of the process are also mapped into the address spaces of other processes and are locked by those processes, the locks established by the other processes are unaffected by a call to the **munlockall** subroutine.

Regarding libraries that are pinned, a distinction has been made internally between a user referencing memory to perform a task related to the application and the system referencing memory on behalf of the application. The former is pinned, and the latter is not. The user-addressable loader data that remains unlocked includes:

- loader entries
- user loader entries
- page-descriptor segment
- usla heap segment
- usla text segment
- all the global segments related to the 64-bit shared library loadlist (shlib heap segment, shlib le segment, shlib text and data heap segments).

This limit affects implementation only, and it does not cause the API to fail.

Upon successful return from a **mlockall** subroutine that specifies **MCL\_CURRENT**, all currently mapped pages of the process' address space are memory-resident and locked. Upon return from the **munlockall** subroutine, all currently mapped pages of the process' address space are unlocked with respect to the process' address space.

The calling process must have the root user authority to use this subroutine.

## Parameters

*flags* Determines whether the pages to be locked are those currently mapped by the address space of the process, those that are mapped in the future, or both.

## Return Values

Upon successful completion, the **mlockall** subroutine returns 0. Otherwise, no additional memory is locked, and the subroutine returns -1 and sets **errno** to indicate the error.

Upon successful completion, the **munlockall** subroutine returns 0. Otherwise, no additional memory is unlocked, and the subroutine returns -1 and sets **errno** to indicate the error.

## Error Codes

The **mlockall** subroutine fails if:

**EINVAL** The *flags* parameter is 0, or includes unimplemented flags or the process has already some plocked memory.

**ENOMEM**

Locking all of the pages currently mapped into the address space of the process would exceed the limit on the amount of memory that the process may lock.

**EPERM**

The calling process does not have the appropriate authority to perform the requested operation.

The **munlockall** subroutine fails if:

**EINVAL**

The process has already some plocked memory

**EPERM**

The calling process does not have the appropriate privilege to perform the requested operation

## Related Information

“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236, “exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243, “fork, f\_fork, or vfork Subroutine” on page 291, “mlock and munlock Subroutine” on page 840, and “munmap Subroutine” on page 893.

---

## mmap or mmap64 Subroutine

### Purpose

Maps a file-system object into virtual memory.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <sys/mman.h>
```

```
void *mmap (addr, len, prot, flags, fildes, off)
void * addr;
size_t len;
int prot, flags, fildes;
off_t off;
```

```
void *mmap64 (addr, len, prot, flags, fildes, off)
void * addr;
size_t len;
int prot, flags, fildes;
off64_t off;
```

### Description

**Attention:** A file-system object should not be simultaneously mapped using both the **mmap** and **shmat** subroutines. Unexpected results may occur when references are made beyond the end of the object.

The **mmap** subroutine creates a new mapped file or anonymous memory region by establishing a mapping between a process-address space and a file-system object. Care needs to be taken when using the **mmap** subroutine if the program attempts to map itself. If the page containing executing instructions is currently referenced as data through an mmap mapping, the program will hang. Use the -H4096 binder option, and that will put the executable text on page boundaries. Then reset the file that contains the executable material, and view via an **mmap** mapping.

A region created by the **mmap** subroutine cannot be used as the buffer for read or write operations that involve a device. Similarly, an **mmap** region cannot be used as the buffer for operations that require either a **pin** or **xmattach** operation on the buffer.

Modifications to a file-system object are seen consistently, whether accessed from a mapped file region or from the **read** or **write** subroutine.

Child processes inherit all mapped regions from the parent process when the **fork** subroutine is called. The child process also inherits the same sharing and protection attributes for these mapped regions. A successful call to any **exec** subroutine will unmap all mapped regions created with the **mmap** subroutine.

The **mmap64** subroutine is identical to the **mmap** subroutine except that the starting offset for the file mapping is specified as a 64-bit value. This permits file mappings which start beyond **OFF\_MAX**.

In the large file enabled programming environment, **mmap** is redefined to be **mmap64**.

If the application has requested SPEC1170 compliant behavior then the **st\_atime** field of the mapped file is marked for update upon successful completion of the **mmap** call.

If the application has requested SPEC1170 compliant behavior then the **st\_ctime** and **st\_mtime** fields of a file that is mapped with **MAP\_SHARED** and **PROT\_WRITE** are marked for update at the next call to **msync** subroutine or **munmap** subroutine if the file has been modified.

## Parameters

<i>addr</i>	Specifies the starting address of the memory region to be mapped. When the <b>MAP_FIXED</b> flag is specified, this address must be a multiple of the page size returned by the <b>sysconf</b> subroutine using the <b>_SC_PAGE_SIZE</b> value for the <i>Name</i> parameter. A region is never placed at address zero, or at an address where it would overlap an existing region.
<i>len</i>	Specifies the length, in bytes, of the memory region to be mapped. The system performs mapping operations over whole pages only. If the <i>len</i> parameter is not a multiple of the page size, the system will include in any mapping operation the address range between the end of the region and the end of the page containing the end of the region.

*prot* Specifies the access permissions for the mapped region. The **sys/mman.h** file defines the following access options:

**PROT\_READ**

Region can be read.

**PROT\_WRITE**

Region can be written.

**PROT\_EXEC**

Region can be executed.

**PROT\_NONE**

Region cannot be accessed.

The *prot* parameter can be the **PROT\_NONE** flag, or any combination of the **PROT\_READ** flag, **PROT\_WRITE** flag, and **PROT\_EXEC** flag logically ORed together. If the **PROT\_NONE** flag is not specified, access permissions may be granted to the region in addition to those explicitly requested. However, write access will not be granted unless the **PROT\_WRITE** flag is specified.

**Note:** The operating system generates a **SIGSEGV** signal if a program attempts an access that exceeds the access permission given to a memory region. For example, if the **PROT\_WRITE** flag is not specified and a program attempts a write access, a **SIGSEGV** signal results.

If the region is a mapped file that was mapped with the **MAP\_SHARED** flag, the **mmap** subroutine grants read or execute access permission only if the file descriptor used to map the file was opened for reading. It grants write access permission only if the file descriptor was opened for writing.

If the region is a mapped file that was mapped with the **MAP\_PRIVATE** flag, the **mmap** subroutine grants read, write, or execute access permission only if the file descriptor used to map the file was opened for reading. If the region is an anonymous memory region, the **mmap** subroutine grants all requested access permissions.

*fildev* Specifies the file descriptor of the file-system object or of the shared memory object to be mapped. If the **MAP\_ANONYMOUS** flag is set, the *fildev* parameter must be -1. After the successful completion of the **mmap** subroutine, the file or the shared memory object specified by the *fildev* parameter can be closed without affecting the mapped region or the contents of the mapped file. Each mapped region creates a file reference, similar to an open file descriptor, which prevents the file data from being deallocated.

**Note:** The **mmap** subroutine supports the mapping of shared memory object and regular files only. An **mmap** call that specifies a file descriptor for a special file fails, returning the **ENODEV** error code. An example of a file descriptor for a special file is one that might be used for mapping either I/O or device memory.

*off* Specifies the file byte offset at which the mapping starts. This offset must be a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.

*flags*

Specifies attributes of the mapped region. Values for the *flags* parameter are constructed by a bitwise-inclusive ORing of values from the following list of symbolic names defined in the **sys/mman.h** file:

#### **MAP\_FILE**

Specifies the creation of a new mapped file region by mapping the file associated with the *fildev* file descriptor. The mapped region can extend beyond the end of the file, both at the time when the **mmap** subroutine is called and while the mapping persists. This situation could occur if a file with no contents was created just before the call to the **mmap** subroutine, or if a file was later truncated. However, references to whole pages following the end of the file result in the delivery of a **SIGBUS** signal. Only one of the **MAP\_FILE** and **MAP\_ANONYMOUS** flags must be specified with the **mmap** subroutine.

#### **MAP\_ANONYMOUS**

Specifies the creation of a new, anonymous memory region that is initialized to all zeros. This memory region can be shared only with the descendants of the current process. When using this flag, the *fildev* parameter must be -1. Only one of the **MAP\_FILE** and **MAP\_ANONYMOUS** flags must be specified with the **mmap** subroutine.

#### **MAP\_VARIABLE**

Specifies that the system select an address for the new memory region if the new memory region cannot be mapped at the address specified by the *addr* parameter, or if the *addr* parameter is null. Only one of the **MAP\_VARIABLE** and **MAP\_FIXED** flags must be specified with the **mmap** subroutine.

#### **MAP\_FIXED**

Specifies that the mapped region be placed exactly at the address specified by the *addr* parameter. If the application has requested SPEC1170 compliant behavior and the **mmap** request is successful, the mapping replaces any previous mappings for the process' pages in the specified range. If the application has not requested SPEC1170 compliant behavior and a previous mapping exists in the range then the request fails. Only one of the **MAP\_VARIABLE** and **MAP\_FIXED** flags must be specified with the **mmap** subroutine.

#### **MAP\_SHARED**

When the **MAP\_SHARED** flag is set, modifications to the mapped memory region will be visible to other processes that have mapped the same region using this flag. If the region is a mapped file region, modifications to the region will be written to the file.

You can specify only one of the **MAP\_SHARED** or **MAP\_PRIVATE** flags with the **mmap** subroutine. **MAP\_PRIVATE** is the default setting when neither flag is specified unless you request SPEC1170 compliant behavior. In this case, you must choose either **MAP\_SHARED** or **MAP\_PRIVATE**.

#### **MAP\_PRIVATE**

When the **MAP\_PRIVATE** flag is specified, modifications to the mapped region by the calling process are not visible to other processes that have mapped the same region. If the region is a mapped file region, modifications to the region are not written to the file.

If this flag is specified, the initial write reference to an object page creates a private copy of that page and redirects the mapping to the copy. Until then, modifications to the page by processes that have mapped the same region with the **MAP\_SHARED** flag are visible.

You can specify only one of the **MAP\_SHARED** or **MAP\_PRIVATE** flags with the **mmap** subroutine. **MAP\_PRIVATE** is the default setting when neither flag is specified unless you request SPEC1170 compliant behavior. In this case, you must choose either **MAP\_SHARED** or **MAP\_PRIVATE**.

## **Return Values**

If successful, the **mmap** subroutine returns the address at which the mapping was placed. Otherwise, it returns -1 and sets the **errno** global variable to indicate the error.

## Error Codes

Under the following conditions, the **mmap** subroutine fails and sets the **errno** global variable to:

<b>EACCES</b>	The file referred to by the <i>fildev</i> parameter is not open for read access, or the file is not open for write access and the <b>PROT_WRITE</b> flag was specified for a <b>MAP_SHARED</b> mapping operation. Or, the file to be mapped has enforced locking enabled and the file is currently locked.
<b>EAGAIN</b>	The <i>fildev</i> parameter refers to a device that has already been mapped.
<b>EBADF</b>	The <i>fildev</i> parameter is not a valid file descriptor, or the <b>MAP_ANONYMOUS</b> flag was set and the <i>fildev</i> parameter is not -1.
<b>EFBIG</b>	The mapping requested extends beyond the maximum file size associated with <i>fildev</i> .
<b>EINVAL</b>	The <i>flags</i> or <i>prot</i> parameter is invalid, or the <i>addr</i> parameter or <i>off</i> parameter is not a multiple of the page size returned by the <b>sysconf</b> subroutine using the <b>_SC_PAGE_SIZE</b> value for the <i>Name</i> parameter.
<b>EINVAL</b>	The application has requested SPEC1170 compliant behavior and the value of <i>flags</i> is invalid (neither <b>MAP_PRIVATE</b> nor <b>MAP_SHARED</b> is set).
<b>EMFILE</b>	The application has requested SPEC1170 compliant behavior and the number of mapped regions would exceed implementation-dependent limit (per process or per system).
<b>ENODEV</b>	The <i>fildev</i> parameter refers to an object that cannot be mapped, such as a terminal.
<b>ENOMEM</b>	There is not enough address space to map <i>len</i> bytes, or the application has not requested Single UNIX Specification, Version 2 compliant behavior and the <b>MAP_FIXED</b> flag was set and part of the address-space range ( <i>addr</i> , <i>addr+len</i> ) is already allocated.
<b>ENXIO</b>	The addresses specified by the range ( <i>off</i> , <i>off+len</i> ) are invalid for the <i>fildev</i> parameter.
<b>EOVERFLOW</b>	The mapping requested extends beyond the offset maximum for the file description associated with <i>fildev</i> .

## Related Information

The **exec** (“**exec**: **execl**, **execle**, **execlp**, **execv**, **execve**, **execvp**, or **exec** Subroutine” on page 236) subroutine, **fork** (“**fork**, **f\_fork**, or **vfork** Subroutine” on page 291) subroutine, **munmap** (“**munmap** Subroutine” on page 893) subroutine, **read** subroutine, **shm\_open** subroutine, **shm\_unlink** subroutine, **shmat** subroutine, **sysconf** subroutine, **write** subroutine.

The **pin** kernel service, **xmattach** kernel service.

List of Memory Manipulation Services, List of Memory Mapping Services, Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## mntctl Subroutine

### Purpose

Returns information about the mount status of the system.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <sys/mntctl.h>
#include <sys/vmount.h>
```

```
int mntctl ( Command, Size, Buffer)
int Command;
int Size;
char *Buffer;
```

## Description

The **mntctl** subroutine is used to query the status of virtual file systems (also known as *mounted* file systems).

Each virtual file system (VFS) is described by a **vmount** structure. This structure is supplied when the VFS is created by the **vmount** subroutine. The **vmount** structure is defined in the **sys/vmount.h** file.

## Parameters

<i>Command</i>	Specifies the operation to be performed. Valid commands are defined in the <b>sys/vmount.h</b> file. At present, the only command is:  <b>MCTL_QUERY</b> Query mount information.
<i>Buffer</i>	Points to a data area that will contain an array of <b>vmount</b> structures. This data area holds the information returned by the query command. Since the <b>vmount</b> structure is variable-length, it is necessary to reference the <i>vmt_length</i> field of each structure to determine where in the <i>Buffer</i> area the next structure begins.
<i>Size</i>	Specifies the length, in bytes, of the buffer pointed to by the <i>Buffer</i> parameter.

## Return Values

If the **mntctl** subroutine is successful, the number of **vmount** structures copied into the *Buffer* parameter is returned. If the *Size* parameter indicates the supplied buffer is too small to hold the **vmount** structures for all the current VFSs, the **mntctl** subroutine sets the first word of the *Buffer* parameter to the required size (in bytes) and returns the value 0. If the **mntctl** subroutine otherwise fails, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **mntctl** subroutine fails and the requested operation is not performed if one or both of the following are true:

<b>EINVAL</b>	The <i>Command</i> parameter is not <b>MCTL_QUERY</b> , or the <i>Size</i> parameter is not a positive value.
<b>EFAULT</b>	The <i>Buffer</i> parameter points to a location outside of the allocated address space of the process.

## Related Information

The **uvmount** or **umount** subroutine, **vmount** or **mount** subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## modf, modff, modfl, modfd32, modfd64, and modfd128 Subroutines

### Purpose

Decomposes a floating-point number.

### Syntax

```
#include <math.h>

float modff (x, iptr)
float x;
float *iptr;

double modf (x, iptr)
```



```

double x, *iptr;

long double modfl (x, iptr)
long double x, *iptr;

_Decimal32 modfd32 (x, iptr)
_Decimal32 x, *iptr;

_Decimal64 modfd64 (x, iptr)
_Decimal64 x, *iptr;

_Decimal128 modf128 (x, iptr)
_Decimal128 x, *iptr;

```

## Description

The **modff**, **modf**, **modfl**, **modfd32**, **modfd64**, and **modfd128** subroutines divide the *x* parameter into integral and fractional parts, each of which has the same sign as the arguments. These subroutines store the integral part as a floating-point value in the object pointed to by the *iptr* parameter.

## Parameters

*x* Specifies the value to be computed.  
*iptr* Points to the object where the integral part is stored.

## Return Values

Upon successful completion, the **modff**, **modf**, **modfl**, **modfd32**, **modfd64**, and **modfd128** subroutines return the signed fractional part of *x*.

If *x* is NaN, a NaN is returned, and *\*iptr* is set to a NaN.

If *x* is  $\pm\text{Inf}$ ,  $\pm 0$  is returned, and *\*iptr* is set to  $\pm\text{Inf}$ .

## Related Information

“class, \_class, finite, isnan, or unordered Subroutines” on page 167 and “ldexp, ldexpf, or ldexpl Subroutine” on page 699.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long Double Floating-Point Format in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## moncontrol Subroutine

### Purpose

Starts and stops execution profiling after initialization by the **monitor** subroutine.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <mon.h>
```

```
int moncontrol ( Mode)  
int Mode;
```

## Description

The **moncontrol** subroutine starts and stops profiling after profiling has been initialized by the **monitor** subroutine. It may be used with either **-p** or **-pg** profiling. When **moncontrol** stops profiling, no output data file is produced. When profiling has been started by the **monitor** subroutine and the **exit** subroutine is called, or when the **monitor** subroutine is called with a value of 0, then profiling is stopped, and an output file is produced, regardless of the state of profiling as set by the **moncontrol** subroutine.

The **moncontrol** subroutine examines global and parameter data in the following order:

1. When the **\_mondata.prof\_type** global variable is neither -1 (**-p** profiling defined) nor +1 (**-pg** profiling defined), no action is performed, 0 is returned, and the function is considered complete.

The global variable is set to -1 in the **mcrt0.o** file and to +1 in the **gcrt0.o** file and defaults to 0 when the **crt0.o** file is used.

2. When the *Mode* parameter is 0, profiling is stopped. For any other value, profiling is started.

The following global variables are used in a call to the **profil** (“profil Subroutine” on page 1176) subroutine:

<b>_mondata.ProfBuf</b>	Buffer address
<b>_mondata.ProfBufSiz</b>	Buffer size/multirange flag
<b>_mondata.ProfLoPC</b>	PC offset for hist buffer - I/O limit
<b>_mondata.ProfScale</b>	PC scale/compute scale flag.

These variables are initialized by the **monitor** subroutine each time it is called to start profiling.

## Parameters

*Mode* Specifies whether to start (resume) or stop profiling.

## Return Values

The **moncontrol** subroutine returns the previous state of profiling. When the previous state was STOPPED, a 0 is returned. When the previous state was STARTED, a 1 is returned.

## Error Codes

When the **moncontrol** subroutine detects an error from the call to the **profil** subroutine, a -1 is returned.

## Related Information

The **monitor** (“monitor Subroutine” on page 851) subroutine, **monstartup** (“monstartup Subroutine” on page 856) subroutine, **profil** (“profil Subroutine” on page 1176) subroutine.

List of Memory Manipulation Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## monitor Subroutine

### Purpose

Starts and stops execution profiling using data areas defined in the function parameters.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mon.h>
```

```
int monitor ( LowProgramCounter, HighProgramCounter, Buffer, BufferSize, NFunction)
```

OR

```
int monitor ( NotZeroA, DoNotCareA, Buffer,-1, NFunction)
```

OR

```
int monitor((caddr_t)0)
```

```
caddr_t LowProgramCounter, HighProgramCounter;
```

```
HISTCOUNTER *Buffer;
```

```
int BufferSize, NFunction;
```

```
caddr_t NotZeroA, DoNotCareA;
```

### Description

The **monitor** subroutine initializes the buffer area and starts profiling, or else stops profiling and writes out the accumulated profiling data. Profiling, when started, causes periodic sampling and recording of the program location within the program address ranges specified. Profiling also accumulates function call count data compiled with the **-p** or **-pg** option.

Executable programs created with the **cc -p** or **cc -pg** command automatically include calls to the **monitor** subroutine (through the **monstartup** and **exit** subroutines) to profile the complete user program, including system libraries. In this case, you do not need to call the **monitor** subroutine.

The **monitor** subroutine is called by the **monstartup** subroutine to begin profiling and by the **exit** subroutine to end profiling. The **monitor** subroutine requires a global data variable to define which kind of profiling, **-p** or **-pg**, is in effect. The **monitor** subroutine initializes four global variables that are used as parameters to the **profil** subroutine by the **moncontrol** subroutine:

- The **monitor** subroutine calls the **moncontrol** subroutine to start the profiling data gathering.
- The **moncontrol** subroutine calls the **profil** subroutine to start the system timer-driven program address sampling.
- The **prof** command processes the data file produced by **-p** profiling.
- The **gprof** command processes the data file produced by **-pg** profiling.

The **monitor** subroutine examines the global data and parameter data in this order:

1. When the **\_monddata.prof\_type** global variable is neither -1 (**-p** profiling defined) nor +1 (**-pg** profiling defined), an error is returned, and the function is considered complete.

The global variable is set to -1 in the **mcrto.o** file and to +1 in the **gcrt0.o** file, and defaults to 0 when the **crt0.o** file is used.

- When the first parameter to the **monitor** subroutine is 0, profiling is stopped and the data file is written out.

If **-p** profiling was in effect, then the file is named **mon.out**. If **-pg** profiling was in effect, the file is named **gmon.out**. The function is complete.

- When the first parameter to the **monitor** subroutine is not , the **monitor** parameters and the profiling global variable, **\_monddata.prof\_type**, are examined to determine how to start profiling.
- When the *BufferSize* parameter is not -1, a single program address range is defined for profiling, and the first **monitor** definition in the syntax is used to define the single program range.
- When the *BufferSize* parameter is -1, multiple program address ranges are defined for profiling, and the second **monitor** definition in the syntax is used to define the multiple ranges. In this case, the *ProfileBuffer* value is the address of an array of **prof** structures. The size of the **prof** array is denoted by a zero value for the *HighProgramCounter* (*p\_high*) field of the last element of the array. Each element in the array, except the last, defines a single programming address range to be profiled. Programming ranges must be in ascending order of the program addresses with ascending order of the **prof** array index. Program ranges may not overlap.

The buffer space defined by the *p\_buff* and *p\_bufsize* fields of all of the **prof** entries must define a single contiguous buffer area. Space for the function-count data is included in the first range buffer. Its size is defined by the *NFunction* parameter. The *p\_scale* entry in the **prof** structure is ignored. The **prof** structure is defined in the **themon.h** file. It contains the following fields:

```
caddr_t p_low;          /* low sampling address */
caddr_t p_high;        /* high sampling address */
HISTCOUNTER *p_buff;   /* address of sampling buffer */
int p_bufsize;        /* buffer size- monitor/HISTCOUNTERs,\
                      profil/bytes */
uint p_scale;         /* scale factor */
```

## Parameters

*LowProgramCounter*  
(**prof** name: *p\_low*)

Defines the lowest execution-time program address in the range to be profiled. The value of the *LowProgramCounter* parameter cannot be 0 when using the **monitor** subroutine to begin profiling.

*HighProgramCounter*  
(**prof** name: *p\_high*)

Defines the next address after the highest-execution time program address in the range to be profiled.

The program address parameters may be defined by function names or address expressions. If defined by a function name, then a function name expression must be used to dereference the function pointer to get the address of the first instruction in the function. This is required because the function reference in this context produces the address of the function descriptor. The first field of the descriptor is the address of the function code. See the examples for typical expressions to use.

*Buffer* (**prof** name: *p\_buff*)

Defines the beginning address of an array of *BufferSize* HISTCOUNTERs to be used for data collection. This buffer includes the space for the program address-sampling counters and the function-count data areas. In the case of a multiple range specification, the space for the function-count data area is included at the beginning of the first range in the *BufferSize* specification.

*BufferSize*  
(**prof** name: p\_bufsize)

Defines the size of the buffer in number of HISTCOUNTERS. Each counter is of type HISTCOUNTER (defined as short in the **mon.h** file). When the buffer includes space for the function-count data area (single range specification and first range of a multi-range specification) the *NFunction* parameter defines the space to be used for the function count data, and the remainder is used for program-address sampling counters for the range defined. The scale for the **profil** call is calculated from the number of counters available for program address-sample counting and the address range defined by the *LowProgramCounter* and *HighProgramCounter* parameters. See the **mon.h** file.

*NFunction*

Defines the size of the space to be used for the function-count data area. The space is included as part of the first (or only) range buffer.

When **-p** profiling is defined, the *NFunction* parameter defines the maximum number of functions to be counted. The space required for each function is defined to be:

```
sizeof(struct poutcnt)
```

The **poutcnt** structure is defined in the **mon.h** file. The total function-count space required is:

```
NFunction * sizeof(struct poutcnt)
```

When **-pg** profiling is defined, the *NFunction* parameter defines the size of the space (in bytes) available for the function-count data structures, as follows:

```
range = HighProgramCounter - LowProgramCounter;  
tonum = TO_NUM_ELEMENTS( range );  
if ( tonum < MINARCS ) tonum = MINARCS;  
if ( tonum > TO_MAX-1 ) tonum = TO_MAX-1;  
tosize = tonum * sizeof( struct tostruct );  
fromsize = FROM_STG_SIZE( range );  
rangesize = tosize + fromsize + sizeof(struct  
gfctl);
```

This is computed and summed for all defined ranges. In this expression, the functions and variables in capital letters as well as the structures are defined in the **mon.h** file.

*NotZeroA*

Specifies a value of parameter 1, which is any value except 0. Ignored when it is not zero.

*DoNotCareA*

Specifies a value of parameter 2, of any value, which is ignored.

## Return Values

The **monitor** subroutine returns 0 upon successful completion.

## Error Codes

If an error is found, the **monitor** subroutine sends an error message to **stderr** and returns -1.

## Examples

1. This example shows how to profile the main load module of a program with **-p** profiling:

```

#include <sys/types.h>
#include <mon.h>
main()
{
extern caddr_t etext; /*system end of main module text symbol*/
extern int start(); /*first function in main program*/
extern struct monglobal _mondata; /*profiling global variables*/
struct desc {
    caddr_t begin; /*initial code address*/
    caddr_t toc; /*table of contents address*/
    caddr_t env; /*environment pointer*/
}; /*function descriptor structure*/
struct desc *fd; /*pointer to function descriptor*/
int rc; /*monitor return code*/
int range; /*program address range for profiling*/
int numfunc; /*number of functions*/
HISTCOUNTER *buffer; /*buffer address*/
int numtics; /*number of program address sample counters*/
int BufferSize; /*total buffer size in numbers of HISTCOUNTERS*/
fd = (struct desc*)start; /*init descriptor pointer to start\
function*/
numfunc = 300; /*arbitrary number for example*/
range = etext - fd->begin; /*compute program address range*/
numtics = NUM_HIST_COUNTERS(range); /*one counter for each 4 byte\
inst*/
BufferSize = numtics + ( numfunc*sizeof (struct poutcnt) \
HIST_COUNTER_SIZE ); /*allocate buffer space*/
buffer = (HISTCOUNTER *) malloc (BufferSize * HIST_COUNTER_SIZE);
if ( buffer == NULL ) /*didn't get space, do error recovery\
here*/
    return(-1);
_mondata.prof_type = _PROF_TYPE_IS_P; /*define -p profiling*/
rc = monitor( fd->begin, (caddr_t)etext, buffer, BufferSize, \
numfunc);
/*start*/
if ( rc != 0 ) /*profiling did not start, do error recovery\
here*/
    return(-1);
/*other code for analysis*/
rc = monitor( (caddr_t)0); /*stop profiling and write data file\
mon.out*/
if ( rc != 0 ) /*did not stop correctly, do error recovery here*/
    return (-1);
}

```

2. This example profiles the main program and the **libc.a** shared library with **-p** profiling. The range of addresses for the shared **libc.a** is assumed to be:

```

low = d0300000
high = d0312244

```

These two values can be determined from the **loadquery** subroutine at execution time, or by using a debugger to view the loaded programs' execution addresses and the loader map.

```

#include <sys/types.h>
#include <mon.h>
main()
{
extern caddr_t etext; /*system end of text symbol*/
extern int start(); /*first function in main program*/
extern struct monglobal _mondata; /*profiling global variables*/
struct prof pb[3]; /*prof array of 3 to define 2 ranges*/
int rc; /*monitor return code*/
int range; /*program address range for profiling*/
int numfunc; /*number of functions to count (max)*/
int numtics; /*number of sample counters*/
int num4fcnt; /*number of HISTCOUNTERS used for fun cnt space*/
int BufferSize1; /*first range BufferSize*/
int BufferSize2; /*second range BufferSize*/

```

```

caddr_t liblo=0xd0300000; /*lib low address (example only)*/
caddr_t libhi=0xd0312244; /*lib high address (example only)*/
numfunc = 400;          /*arbitrary number for example*/
/*compute first range buffer size*/
range = etext - *(uint *) start; /*init range*/
numtics = NUM_HIST_COUNTERS( range );
/*one counter for each 4 byte inst*/
num4fcnt = numfunc*sizeof( struct poutcnt )/HIST_COUNTER_SIZE;
BufferSize1 = numtics + num4fcnt;
/*compute second range buffer size*/
range = libhi-liblo;
BufferSize2 = range / 12; /*counter for every 12 inst bytes for\
 a change*/
/*allocate buffer space - note: must be single contiguous\
 buffer*/
pb[0].p_buff = (HISTCOUNTER *)malloc( (BufferSize1 +BufferSize2)\
 *HIST_COUNTER_SIZE);
if ( pb[0].p_buff == NULL ) /*didn't get space - do error\
 recovery here*/
    return(-1);
/*set up the first range values*/
pb[0].p_low = *(uint*)start; /*start of main module*/
pb[0].p_high = (caddr_t)etext; /*end of main module*/
pb[0].p_BufferSize = BufferSize1; /*prog addr cnt space + \
 func cnt space*/
/*set up last element marker*/
pb[2].p_high = (caddr_t)0;
_mondata.prof_type = _PROF_TYPE_IS_P; /*define -p\
 profiling*/
rc = monitor( (caddr_t)1, (caddr_t)1, pb, -1, numfunc); \
 /*start*/
if ( rc != 0 ) /*profiling did not start - do error recovery\
 here*/
    return (-1);
/*other code for analysis ...*/
rc = monitor( (caddr_t)0); /*stop profiling and write data \
 file mon.out*/
if ( rc != 0 ) /*did not stop correctly - do error recovery\
 here*/
    return (-1);

```

3. This example shows how to profile contiguously loaded functions beginning at zit up to but not including zot with **-pg** profiling:

```

#include <sys/types.h>
#include <mon.h>
main()
{
extern zit();          /*first function to profile*/
extern zot();          /*upper bound function*/
extern struct monglobal _mondata; /*profiling global variables*/
int rc;                /*monstartup return code*/
_mondata.prof_type = _PROF_TYPE_IS_PG; /*define -pg profiling*/
/*Note cast used to obtain function code addresses*/
rc = monstartup(*(uint *)zit,*(uint *)zot); /*start*/
if ( rc != 0 ) /*profiling did not start, do error recovery\
 here*/
    return(-1);
/*other code for analysis ...*/
exit(0); /*stop profiling and write data file gmon.out*/
}

```

## Files

<b>mon.out</b>	Data file for <b>-p</b> profiling.
<b>gmon.out</b>	Data file for <b>-pg</b> profiling.

`/usr/include/mon.h`

Defines the `_mondata.prof_type` global variable in the `monglobal` data structure, the `prof` structure, and the functions referred to in the previous examples.

## Related Information

The `moncontrol` (“moncontrol Subroutine” on page 849) subroutine, `monstartup` (“monstartup Subroutine”) subroutine, `profil` (“profil Subroutine” on page 1176) subroutine.

The `gprof` command, `prof` command.

The `_end,_etext`, or `_edata` (“\_end, \_etext, or \_edata Identifier” on page 224) Identifier.

List of Memory Manipulation Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## monstartup Subroutine

### Purpose

Starts and stops execution profiling using default-sized data areas.

### Library

Standard C Library (`libc.a`)

### Syntax

```
#include <mon.h>
```

```
int monstartup ( LowProgramCounter, HighProgramCounter)
```

OR

```
int monstartup((caddr_t)-1), (caddr_t) FragBuffer)
```

OR

```
int monstartup((caddr_t)-1, (caddr_t)0)
```

```
caddr_t LowProgramCounter;  
caddr_t HighProgramCounter;
```

### Description

The `monstartup` subroutine allocates data areas of default size and starts profiling. Profiling causes periodic sampling and recording of the program location within the program address ranges specified, and accumulation of function-call count data for functions that have been compiled with the `-p` or `-pg` option.

Executable programs created with the `cc -p` or `cc -pg` command automatically include a call to the `monstartup` subroutine to profile the complete user program, including system libraries. In this case, you do not need to call the `monstartup` subroutine.

The `monstartup` subroutine is called by the `mcrt0.o (-p)` file or the `gcrt0.o (-pg)` file to begin profiling. The `monstartup` subroutine requires a global data variable to define whether `-p` or `-pg` profiling is to be in effect. The `monstartup` subroutine calls the `monitor` subroutine to initialize the data areas and start profiling.



The **prof** command is used to process the data file produced by **-p** profiling. The **gprof** command is used to process the data file produced by **-pg** profiling.

The **monstartup** subroutine examines the global and parameter data in the following order:

1. When the **\_mondata.prof\_type** global variable is neither -1 (**-p** profiling defined) nor +1 (**-pg** profiling defined), an error is returned and the function is considered complete.  
The global variable is set to -1 in the **mcrto.o** file and to +1 in the **gcrt0.o** file, and defaults to 0 when **crt0.o** is used.
2. When the *LowProgramCounter* value is not -1:
  - A single program address range is defined for profilingAND
  - The first **monstartup** definition in the syntax is used to define the program range.
3. When the *LowProgramCounter* value is -1 and the *HighProgramCounter* value is not 0:
  - Multiple program address ranges are defined for profilingAND
  - The second **monstartup** definition in the syntax is used to define multiple ranges. The *HighProgramCounter* parameter, in this case, is the address of a **frag** structure array. The **frag** array size is denoted by a zero value for the *HighProgramCounter* (*p\_high*) field of the last element of the array. Each array element except the last defines one programming address range to be profiled. Programming ranges must be in ascending order of the program addresses with ascending order of the **prof** array index. Program ranges may not overlap.
4. When the *LowProgramCounter* value is -1 and the *HighProgramCounter* value is 0:
  - The whole program is defined for profilingAND
  - The third **monstartup** definition in the syntax is used. The program ranges are determined by **monstartup** and may be single range or multirange.

## Parameters

*LowProgramCounter* (**frag** name: *p\_low*)

Defines the lowest execution-time program address in the range to be profiled.

*HighProgramCounter*(**frag** name: *p\_high*)

Defines the next address after the highest execution-time program address in the range to be profiled.

The program address parameters may be defined by function names or address expressions. If defined by a function name, then a function name expression must be used to dereference the function pointer to get the address of the first instruction in the function. This is required because the function reference in this context produces the address of the function descriptor. The first field of the descriptor is the address of the function code. See the examples for typical expressions to use.

*FragBuffer*

Specifies the address of a frag structure array.

## Examples

1. This example shows how to profile the main load module of a program with **-p** profiling:

```
#include <sys/types.h>
#include <mon.h>
main()
{
extern caddr_t etext;      /*system end of text
```

```

symbol*/
extern int start();      /*first function in main\
                        program*/
extern struct monglobal _mondata; /*profiling global variables*/
struct desc {          /*function
descriptor fields*/
    caddr_t begin;     /*initial code
address*/
    caddr_t toc;       /*table of contents
address*/
    caddr_t env;       /*environment
pointer*/
}
;                      /*function
descriptor structure*/
struct desc *fd;       /*pointer to function\
descriptor*/
int rc;                /*monstartup
return code*/
fd = (struct desc *)start; /*init descriptor pointer to\
start
function*/
_mondata.prof_type = _PROF_TYPE_IS_P; /*define -p profiling*/
rc = monstartup( fd->begin, (caddr_t) &text); /*start*/
if ( rc != 0 )         /*profiling did
not start - do\
error
recovery here*/ return(-1);
/*other code
for analysis ...*/
return(0);             /*stop profiling and
write data\
file
mon.out*/
}

```

2. This example shows how to profile the complete program with **-p** profiling:

```

#include <sys/types.h>
#include <mon.h>
main()
{
extern struct monglobal _mondata; /*profiling global\
variables*/
int rc;                /*monstartup
return code*/
_mondata.prof_type = _PROF_TYPE_IS_P; /*define -p profiling*/
rc = monstartup( (caddr_t)-1, (caddr_t)0); /*start*/
if ( rc != 0 )         /*profiling did
not start -\

do error recovery here*/
return (-1);
/*other code
for analysis ...*/
return(0);             /*stop profiling and
write data\
file
mon.out*/
}

```

3. This example shows how to profile contiguously loaded functions beginning at `zit` up to but not including `zot` with **-pg** profiling:

```

#include <sys/types.h>
#include <mon.h>
main()
{
extern zit();          /*first function

```

```

to profile*/
extern zot();          /*upper bound
function*/
extern struct monglobal _mondata; /*profiling global variables*/
int rc;                /*monstartup
return code*/
_mondata.prof_type = _PROF_TYPE_IS_PG; /*define -pg profiling*/
/*Note cast used to obtain function code addresses*/
rc = monstartup(*(uint *)zit,*(uint *)zot); /*start*/
if ( rc != 0 )        /*profiling did
not start - do\
                error
recovery here*/
return(-1);
                /*other code
for analysis ...*/
exit(0); /*stop profiling and write data file gmon.out*/
}

```

## Return Values

The **monstartup** subroutine returns 0 upon successful completion.

## Error Codes

If an error is found, the **monstartup** subroutine outputs an error message to **stderr** and returns -1.

## Files

<b>mon.out</b>	Data file for <b>-p</b> profiling.
<b>gmon.out</b>	Data file for <b>-pg</b> profiling.
<b>mon.h</b>	Defines the <b>_mondata.prof_type</b> variable in the <b>monglobal</b> data structure, the <b>prof</b> structure, and the functions referred to in the examples.

## Related Information

The **moncontrol** (“moncontrol Subroutine” on page 849) subroutine, **monitor** (“monitor Subroutine” on page 851) subroutine, **profil** (“profil Subroutine” on page 1176) subroutine.

The **gprof** command, **prof** command.

The **\_end**, **\_etext**, or **\_edata** (“\_end, \_etext, or \_edata Identifier” on page 224) Identifier.

List of Memory Manipulation Services in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## mprotect Subroutine

### Purpose

Modifies access protections for memory mapping or shared memory.

### Library

Standard C Library (**libc.a**)

### Syntax

```

#include <sys/types.h>
#include <sys/mman.h>

```

```
int mprotect ( addr, len, prot)
void *addr;
size_t len;
int prot;
```

## Description

The **mprotect** subroutine modifies the access protection of a mapped file or shared memory region or anonymous memory region created by the **mmap** subroutine. Processes running in an environment where the MPROTECT\_SHM=ON environmental variable is defined can also use the **mprotect** subroutine to modify the access protection of a shared memory region created by the **shmget**, **ra\_shmget**, or **ra\_shmgetv** subroutine and attached by the **shmat** subroutine.

The user who protects shared memory with the **mprotect** subroutine must be either the user who created the shared memory descriptor, the user who owns the shared memory descriptor, or the root user.

The **mprotect** subroutine can only be used on shared memory regions backed with 4 KB or 64 KB pages; shared memory regions backed by 16 MB and 16 GB pages are not supported by the **mprotect** subroutine. The page size used to back a shared memory region can be obtained using the **vmgetinfo** subroutine and specifying VM\_PAGE\_INFO for the *command* parameter.

The **mprotect** subroutine cannot be used for shared memory that has been pre-translated. This includes shared memory regions created with the SHM\_PIN flag specified to the **shmget** subroutine as well as shared memory regions that have been pinned using the **shmctl** subroutine with the SHM\_LOCK flag specified.

## Parameters

*addr* Specifies the address of the region to be modified. Must be a multiple of the page size backing the memory region.

*len* Specifies the length, in bytes, of the region to be modified. For shared memory regions backed with 4 KB pages, the *len* parameter will be rounded off to the next multiple of the page size. Otherwise, the *len* parameter must be a multiple of the page size backing the memory region.

*prot* Specifies the new access permissions for the mapped region. Legitimate values for the *prot* parameter are the same as those permitted for the **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine, as follows:

### **PROT\_READ**

Region can be read.

### **PROT\_WRITE**

Region can be written.

### **PROT\_EXEC**

Region can be executed.

### **PROT\_NONE**

Region cannot be accessed. PROT\_NONE is not a valid *prot* parameter for shared memory attached with the **shmat** subroutine.

## Return Values

When successful, the **mprotect** subroutine returns 0. Otherwise, it returns -1 and sets the **errno** global variable to indicate the error.

**Note:** The return value for the **mprotect** subroutine is 0 if it fails because the region given was not created by **mmap** unless XPG 1170 behavior is requested by setting the **XPG\_SUS\_ENV** environment variable to **ON**.

## Error Codes

If the **mprotect** subroutine is unsuccessful, the **errno** global variable might be set to one of the following values:

**Attention:** If the **mprotect** subroutine is unsuccessful because of a condition other than that specified by the **EINVAL** error code, the access protection for some pages in the (*addr*, *addr + len*) range might have been changed.

<b>EACCES</b>	The <i>prot</i> parameter specifies a protection that conflicts with the access permission set for the underlying file.
<b>EPERM</b>	The user is not the creator or owner of the shared memory region and is not the root user.
<b>ENOTSUP</b>	The <i>prot</i> parameter specified is not valid for the region specified.
<b>EINVAL</b>	The <i>addr</i> or <i>len</i> parameter is not a multiple of the page size backing the memory region.
<b>ENOMEM</b>	The application has requested Single UNIX Specification, Version 2 compliant behavior, but addresses in the range are not valid for the address space of the process, or the addresses specify one or more pages that are not attached to the user's address space by a previous <b>mmap</b> or <b>shmat</b> subroutine call.
<b>ENOTSUP</b>	The shared memory region specified is backed by 64 KB pages, but the <i>addr</i> or <i>len</i> parameter is not 64 KB aligned, or <b>PROT_NONE</b> protection was specified for a shared memory region, or a pre-translated shared memory region was specified, or a shared memory region backed by 16 MB or 16 GB pages was specified.

## Related Information

The **vmgetinfo** subroutine, **shmget** subroutine, **shmctl** subroutine.

---

## mq\_close Subroutine

### Purpose

Closes a message queue.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>

int mq_close (mqdes)
mqd_t mqdes;
```

### Description

The **mq\_close** subroutine removes the association between the message queue descriptor, *mqdes*, and its message queue. The results of using this message queue descriptor after successful return from the **mq\_close** subroutine, and until the return of this message queue descriptor from a subsequent **mq\_open** call, are undefined.

If the process has successfully attached a notification request to the message queue through the *mqdes* parameter, this attachment is removed, and the message queue is available for another process to attach for notification.

## Parameters

*mqdes* Specifies the message queue descriptor.

## Return Values

Upon successful completion, the **mq\_close** subroutine returns a zero. Otherwise, the subroutine returns a -1 and sets **errno** to indicate the error.

## Error Codes

The **mq\_close** subroutine fails if:

<b>EBADF</b>	The <i>mqdes</i> parameter is not a valid message queue descriptor.
<b>ENOMEM</b>	Insufficient memory for the required operation.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_open Subroutine” on page 864 and “mq\_unlink Subroutine” on page 874.

---

## mq\_getattr Subroutine

### Purpose

Gets message queue attributes.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>
```

```
int mq_getattr (mqdes, mqstat)
mqd_t mqdes;
struct mq_attr *mqstat;
```

### Description

The **mq\_getattr** subroutine obtains status information and attributes of the message queue and the open message queue description associated with the message queue descriptor.

The results are returned in the **mq\_attr** structure referenced by the *mqstat* parameter.

Upon return, the following members have the values associated with the open message queue description as set when the message queue was opened and as modified by subsequent calls to the **mq\_setattr** subroutine:

- *mq\_flags*

The following attributes of the message queue are returned as set at message queue creation:

- *mq\_maxmsg*
- *mq\_msgsize*

Upon return, the following member within the **mq\_attr** structure referenced by the *mqstat* parameter is set to the current state of the message queue:

*mq\_curmsgs*                    The number of messages currently on the queue.

## Parameters

*mqdes*                        Specifies a message queue descriptor.  
*mqstat*                      Points to the **mq\_attr** structure.

## Return Values

Upon successful completion, the **mq\_getattr** subroutine returns zero. Otherwise, the subroutine returns -1 and sets **errno** to indicate the error.

## Error Codes

The **mq\_getattr** subroutine fails if:

**EBADF**                      The *mqdes* parameter is not a valid message queue descriptor.  
**EFAULT**                    Invalid user address.  
**EINVAL**                    The *mqstat* parameter value is not valid.  
**ENOMEM**                  Insufficient memory for the required operation.  
**ENOTSUP**                 This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_open Subroutine” on page 864 and “mq\_setattr Subroutine” on page 869.

---

## mq\_notify Subroutine

### Purpose

Notifies a process that a message is available.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>

int mq_notify (mqdes, notification)
mqd_t mqdes;
const struct sigevent *notification;
```

### Description

If the *notification* parameter is not NULL, the **mq\_notify** subroutine registers the calling process to be notified of message arrival at an empty message queue associated with the specified message queue descriptor, *mqdes*. The notification specified by the *notification* parameter is sent to the process when the message queue transitions from empty to non-empty. At any time only one process may be registered for notification by a message queue. If the calling process or any other process has already registered for notification of message arrival at the specified message queue, subsequent attempts to register for that message queue fails.

If notification is NULL and the process is currently registered for notification by the specified message queue, the existing registration is removed.

When the notification is sent to the registered process, its registration is removed. The message queue is then available for registration.

If a process has registered for notification of message arrival at a message queue and a thread is blocked in the **mq\_receive** subroutine waiting to receive a message, the arriving message satisfies the appropriate **mq\_receive**. The resulting behavior is as if the message queue remains empty, and no notification is sent.

## Parameters

*mqdes* Specifies a message queue descriptor.  
*notification* Points to the **sigevent** structure.

## Return Values

Upon successful completion, the **mq\_notify** subroutine returns a zero. Otherwise, it returns a value of -1 and sets **errno** to indicate the error.

## Error Codes

The **mq\_notify** subroutine fails if:

<b>EBADF</b>	The <i>mqdes</i> parameter is not a valid message queue descriptor.
<b>EBUSY</b>	A process is already registered for notification by the message queue.
<b>EFAULT</b>	Invalid used address.
<b>ENOMEM</b>	Insufficient memory for the required operation.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.
<b>EINVAL</b>	The current process is not registered for notification for the specified message queue and registration removal was requested.

## Related Information

“mq\_open Subroutine.”

---

## mq\_open Subroutine

### Purpose

Opens a message queue.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>

mqd_t mq_open (name, oflag [mode, attr])
const char *name;
int oflag;
mode_t mode;
mq_attr *attr;
```

### Description

The **mq\_open** subroutine establishes a connection between a process and a message queue with a message queue descriptor. It creates an open message queue description that refers to the message



queue, and a message queue descriptor that refers to that open message queue description. The message queue descriptor is used by other subroutines to refer to that message queue.

The *name* parameter points to a string naming a message queue, and has no representation in the file system. The *name* parameter conforms to the construction rules for a pathname. It may or may not begin with a slash character, but contains at least one character. Processes calling the **mq\_open** subroutine with the same value of *name* refer to the same message queue object, as long as that name has not been removed. If the *name* parameter is not the name of an existing message queue and creation is not requested, the **mq\_open** subroutine will fail and return an error.

The *oflag* parameter requests the desired receive and send access to the message queue. The requested access permission to receive messages or send messages is granted if the calling process would be granted read or write access, respectively, to an equivalently protected file.

The value of the *oflag* parameter is the bitwise-inclusive OR of values from the following list. Applications specify exactly one of the first three values (access modes) below in the value of the *oflag* parameter:

#### **O\_RDONLY**

Open the message queue for receiving messages. The process can use the returned message queue descriptor with the **mq\_receive** subroutine, but not the **mq\_send** subroutine. A message queue may be open multiple times in the same or different processes for receiving messages.

#### **O\_WRONLY**

Open the queue for sending messages. The process can use the returned message queue descriptor with the **mq\_send** subroutine but not the **mq\_receive** subroutine. A message queue may be open multiple times in the same or different processes for sending messages.

#### **O\_RDWR**

Open the queue for both receiving and sending messages. The process can use any of the functions allowed for the **O\_RDONLY** and **O\_WRONLY** flags. A message queue may be open multiple times in the same or different processes for sending messages.

Any combination of the remaining flags may be specified in the value of the *oflag* parameter:

#### **O\_CREAT**

Create a message queue. It requires two additional arguments: *mode*, which is of **mode\_t** type, and *attr*, which is a pointer to an **mq\_attr** structure. If the pathname *name* has already been used to create a message queue that still exists, this flag has no effect, except as noted under the **O\_EXCL** flag. Otherwise, a message queue is created without any messages in it. The user ID of the message queue is set to the effective user ID of the process, and the group ID of the message queue is set to the effective group ID of the process. The file permission bits are set to the value of *mode*. When bits in the *mode* parameter other than file permission bits are set, they have no effect. If *attr* is NULL, the message queue is created with default message queue attributes. Default values are 128 for *mq\_maxmsg* and 1024 for *mq\_msgsize*. If *attr* is non-NULL, the message queue *mq\_maxmsg* and *mq\_msgsize* attributes are set to the values of the corresponding members in the **mq\_attr** structure referred to by *attr*.

#### **O\_EXCL**

If the **O\_EXCL** and **O\_CREAT** flags are set, the **mq\_open** subroutine fails if the message queue name exists. The check for the existence of the message queue and the creation of the message queue if it does not exist is atomic with respect to other threads executing **mq\_open** naming the same name with the **O\_EXCL** and **O\_CREAT** flags set. If the **O\_EXCL** flag is set and the **O\_CREAT** flag is not set, the **O\_EXCL** flag is ignored.

#### **O\_NONBLOCK**

Determines whether the **mq\_send** or **mq\_receive** subroutine waits for resources or messages that are not currently available, or fails with **errno** set to **EAGAIN**; see “mq\_send Subroutine” on page 868 and “mq\_receive Subroutine” on page 866 for details.

The `mq_open` subroutine does not add or remove messages from the queue.

## Parameters

<i>name</i>	Points to a string naming a message queue.
<i>oflag</i>	Requests the desired receive and send access to the message queue.
<i>mode</i>	Specifies the value of the file permission bits. Used with <b>O_CREAT</b> to create a message queue.
<i>attr</i>	Points to an <b>mq_attr</b> structure. Used with <b>O_CREAT</b> to create a message queue.

## Return Values

Upon successful completion, the `mq_open` subroutine returns a message queue descriptor. Otherwise, it returns `(mqd_t)-1` and sets **errno** to indicate the error.

## Error Codes

The `mq_open` subroutine fails if:

<b>EACCES</b>	The message queue exists and the permissions specified by the <i>oflag</i> parameter are denied.
<b>EEXIST</b>	The <b>O_CREAT</b> and <b>O_EXCL</b> flags are set and the named message queue already exists.
<b>EFAULT</b>	Invalid used address.
<b>EINVAL</b>	The <code>mq_open</code> subroutine is not supported for the given name.
<b>EINVAL</b>	The <b>O_CREAT</b> flag was specified in the <i>oflag</i> parameter, the value of <i>attr</i> is not NULL, and either <i>mq_maxmsg</i> or <i>mq_msgsize</i> was less than or equal to zero.
<b>EINVAL</b>	The <i>oflag</i> parameter value is not valid.
<b>EMFILE</b>	Too many message queue descriptors are currently in use by this process.
<b>ENAMETOOLONG</b>	The length of the <i>name</i> parameter exceeds <b>PATH_MAX</b> or a pathname component is longer than <b>NAME_MAX</b> .
<b>ENFILE</b>	Too many message queues are currently open in the system.
<b>ENOENT</b>	The <b>O_CREAT</b> flag is not set and the named message queue does not exist.
<b>ENOMEM</b>	Insufficient memory for the required operation.
<b>ENOSPC</b>	There is insufficient space for the creation of the new message queue.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_close Subroutine” on page 861, “mq\_getattr Subroutine” on page 862, “mq\_receive Subroutine,” “mq\_send Subroutine” on page 868, “mq\_setattr Subroutine” on page 869, “mq\_unlink Subroutine” on page 874, “msgctl Subroutine” on page 879, “msgget Subroutine” on page 881, “msgrcv Subroutine” on page 882, and “msgsnd Subroutine” on page 885.

---

## mq\_receive Subroutine

### Purpose

Receives a message from a message queue.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>
```

```
ssize_t mq_receive (mqdes, msg_ptr, msg_len, msg_prio)
```

```
mqd_t mqdes;
char *msg_ptr;
size_t msg_len;
unsigned *msg_prio;
```

## Description

The **mq\_receive** subroutine receives the oldest of the highest priority messages from the message queue specified by the *mqdes* parameter. If the size of the buffer in bytes, specified by the *msg\_len* parameter, is less than the *mq\_msgsize* attribute of the message queue, the subroutine fails and returns an error. Otherwise, the selected message is removed from the queue and copied to the buffer pointed to by the *msg\_ptr* parameter.

If the *msg\_prio* parameter is not NULL, the priority of the selected message is stored in the location referenced by *msg\_prio*.

If the specified message queue is empty and the **O\_NONBLOCK** flag is not set in the message queue description associated with the *mqdes* parameter, the **mq\_receive** subroutine blocks until a message is enqueued on the message queue or until **mq\_receive** is interrupted by a signal. If more than one thread is waiting to receive a message when a message arrives at an empty queue and the Priority Scheduling option is supported, the thread of highest priority that has been waiting the longest is selected to receive the message. If the specified message queue is empty and the **O\_NONBLOCK** flag is set in the message queue description associated with the *mqdes* parameter, no message is removed from the queue, and the **mq\_receive** subroutine returns an error.

## Parameters

<i>mqdes</i>	Specifies the message queue descriptor.
<i>msg_ptr</i>	Points to the buffer where the message is copied.
<i>msg_len</i>	Specifies the length of the message, in bytes.
<i>msg_prio</i>	Stores the priority of the selected message.

## Return Values

Upon successful completion, the **mq\_receive** subroutine returns the length of the selected message in bytes and the message is removed from the queue. Otherwise, no message is removed from the queue, and the subroutine returns -1 and sets **errno** to indicate the error.

## Error Codes

The **mq\_receive** subroutine fails if:

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> flag was set in the message description associated with the <i>mqdes</i> parameter, and the specified message queue is empty.
<b>EBADF</b>	The <i>mqdes</i> parameter is not a valid message queue descriptor open for reading.
<b>EFAULT</b>	Invalid used address.
<b>EIDRM</b>	The specified message queue was removed during the required operation.
<b>EINTR</b>	The <b>mq_receive</b> subroutine was interrupted by a signal.
<b>EINVAL</b>	The <i>msg_ptr</i> parameter is null.
<b>EMSGSIZE</b>	The specified message buffer size, <i>msg_len</i> , is less than the message size attribute of the message queue.
<b>ENOMEM</b>	Insufficient memory for the required operation.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_open Subroutine” on page 864 and “mq\_send Subroutine” on page 868.

---

## mq\_send Subroutine

### Purpose

Sends a message to a message queue.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>

int mq_send (mqdes, msg_ptr, msg_len, msg_prio)
mqd_t mqdes;
const char *msg_ptr;
size_t msg_len;
unsigned *msg_prio;
```

### Description

The **mq\_send** subroutine adds the message pointed to by the *msg\_ptr* parameter to the message queue specified by the *mqdes* parameter. The *msg\_len* parameter specifies the length of the message, in bytes, pointed to by *msg\_ptr*. The value of *msg\_len* is less than or equal to the *mq\_msgsize* attribute of the message queue, or the **mq\_send** subroutine will fail.

If the specified message queue is not full, the **mq\_send** subroutine behaves as if the message is inserted into the message queue at the position indicated by the *msg\_prio* parameter. A message with a larger numeric value of *msg\_prio* will be inserted before messages with lower values of *msg\_prio*. A message will be inserted after other messages in the queue with equal *msg\_prio*. The value of *msg\_prio* will be less than **MQ\_PRIO\_MAX**.

If the specified message queue is full and **O\_NONBLOCK** is not set in the message queue description associated with *mqdes*, the **mq\_send** subroutine will block until space becomes available to enqueue the message, or until **mq\_send** is interrupted by a signal. If more than one thread is waiting to send when space becomes available in the message queue and the Priority Scheduling option is supported, the thread of the highest priority that has been waiting the longest is unblocked to send its message. Otherwise, it is unspecified which waiting thread is unblocked. If the specified message queue is full and **O\_NONBLOCK** is set in the message queue description associated with *mqdes*, the message is not queued and the **mq\_send** subroutine returns an error.

### Parameters

<i>mqdes</i>	Specifies the message queue descriptor.
<i>msg_ptr</i>	Points to the message to be added.
<i>msg_len</i>	Specifies the length of the message, in bytes.
<i>msg_prio</i>	Specifies the position of the message in the message queue.

### Return Values

Upon successful completion, the **mq\_send** subroutine returns a zero. Otherwise, no message is enqueued, the subroutine returns -1, and **errno** is set to indicate the error.

## Error Codes

The `mq_send` subroutine fails if:

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> flag is set in the message queue description associated with the <i>mqdes</i> parameter, and the specified message queue is full (maximum number of messages in the queue or maximum number of bytes in the queue is reached).
<b>EBADF</b>	The <i>mqdes</i> parameter is not a valid message queue descriptor open for writing.
<b>EFAULT</b>	Invalid used address.
<b>EIDRM</b>	The specified message queue was removed during the required operation.
<b>EINTR</b>	A signal interrupted the call to the <code>mq_send</code> subroutine.
<b>EINVAL</b>	The value of the <i>msg_prio</i> parameter was outside the valid range.
<b>EINVAL</b>	The <i>msg_ptr</i> parameter is null.
<b>EMSGSIZE</b>	The specified message length, <i>msg_len</i> , exceeds the message size attribute of the message queue.
<b>ENOMEM</b>	Insufficient memory for the required operation.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_open Subroutine” on page 864 and “mq\_receive Subroutine” on page 866.

---

## mq\_setattr Subroutine

### Purpose

Sets message queue attributes.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqqueue.h>

int mq_setattr (mqdes, mqstat, omqstat)
mqd_t mqdes;
const struct mq_attr *mqstat;
struct mq_attr *omqstat;
```

### Description

The `mq_setattr` subroutine sets attributes associated with the open message queue description referenced by the message queue descriptor specified by *mqdes*.

The message queue attributes corresponding to the following members defined in the **mq\_attr** structure are set to the specified values upon successful completion of the `mq_setattr` subroutine.

The value of the *mq\_flags* member is either zero or **O\_NONBLOCK**.

The values of the *mq\_maxmsg*, *mq\_msgsize*, and *mq\_curmsgs* members of the **mq\_attr** structure are ignored by the `mq_setattr` subroutine.

If the *omqstat* parameter is non-NULL, the `mq_setattr` subroutine stores, in the location referenced by *omqstat*, the previous message queue attributes and the current queue status. These values are the same as would be returned by a call to the `mq_getattr` subroutine at that point.

## Parameters

<i>mqdes</i>	Specifies the message queue descriptor.
<i>mqstat</i>	Specifies the status of the message queue.
<i>omqstat</i>	Specifies the status of the previous message queue.

## Return Values

Upon successful completion, the **mq\_setattr** subroutine returns a zero and the attributes of the message queue are changed as specified.

Otherwise, the message queue attributes are unchanged, and the subroutine returns a -1 and sets **errno** to indicate the error.

## Error Codes

The **mq\_setattr** subroutine fails if:

<b>EBADF</b>	The <i>mqdes</i> parameter is not a valid message queue descriptor.
<b>EFAULT</b>	Invalid user address.
<b>EINVAL</b>	The <i>mqstat</i> parameter value is not valid.
<b>ENOMEM</b>	Insufficient memory for the required operation.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_open Subroutine” on page 864 and “mq\_getattr Subroutine” on page 862.

---

## mq\_receive, mq\_timedreceive Subroutine

### Purpose

Receives a message from a message queue (REALTIME).

### Syntax

```
#include <mqueue.h>

ssize_t mq_receive(mqd_t mqdes, char *msg_ptr,
                  size_t msg_len, unsigned *msg_prio,

#include <mqueue.h>
#include <time.h>

ssize_t mq_timedreceive(mqd_t mqdes, char *restrict msg_ptr,
                       size_t msg_len, unsigned *restrict msg_prio,
                       const struct timespec *restrict abs_timeout);
```

### Description

The **mq\_receive()** function receives the oldest of the highest priority messages from the message queue specified by *mqdes*. If the size of the buffer, in bytes, specified by the *msg\_len* argument is less than the *mq\_msgsize* attribute of the message queue, the function fails and returns an error. Otherwise, the selected message is removed from the queue and copied to the buffer pointed to by the *msg\_ptr* argument.

If the value of *msg\_len* is greater than {SSIZE\_MAX}, the result is implementation-defined.

If the *msg\_prio* argument is not NULL, the priority of the selected message is stored in the location referenced by *msg\_prio*.

If the specified message queue is empty and O\_NONBLOCK is not set in the message queue description associated with *mqdes*, **mq\_receive()** blocks until a message is enqueued on the message queue or until **mq\_receive()** is interrupted by a signal. If more than one thread is waiting to receive a message when a message arrives at an empty queue and the Priority Scheduling option is supported, then the thread of highest priority that has been waiting the longest is selected to receive the message. Otherwise, it is unspecified which waiting thread receives the message. If the specified message queue is empty and O\_NONBLOCK is set in the message queue description associated with *mqdes*, no message is removed from the queue, and **mq\_receive()** returns an error.

The **mq\_timedreceive()** function receives the oldest of the highest priority messages from the message queue specified by *mqdes* as described for the **mq\_receive()** function. However, if O\_NONBLOCK was not specified when the message queue was opened by the **mq\_open()** function, and no message exists on the queue to satisfy the receive, the wait for such a message is terminated when the specified timeout expires. If O\_NONBLOCK is set, this function matches **mq\_receive()**.

The timeout expires when the absolute time specified by *abs\_timeout* passes—as measured by the clock on which timeouts are based (that is, when the value of that clock equals or exceeds *abs\_timeout*), or when the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

If the **Timers** option is supported, the timeout is based on the CLOCK\_REALTIME clock; if the **Timers** option is not supported, the timeout is based on the system clock as returned by the **time()** function.

The resolution of the timeout matches the resolution of the clock on which it is based. The *timespec* argument is defined in the `<time.h>` header.

The operation never fails with a timeout if a message can be removed from the message queue immediately. The validity of the *abs\_timeout* parameter does not need to be checked if a message can be removed from the message queue immediately.

## Return Values

Upon successful completion, the **mq\_receive()** and **mq\_timedreceive()** functions return the length of the selected message in bytes and the message is removed from the queue. Otherwise, no message shall be removed from the queue, the functions return a value of -1, and *errno* is set to indicate the error.

## Error Codes

The **mq\_receive()** and **mq\_timedreceive()** functions fail if:

[EAGAIN]	O_NONBLOCK was set in the message description associated with <i>mqdes</i> , and the specified message queue is empty.
[EBADF]	The <i>mqdes</i> argument is not a valid message queue descriptor open for reading.
[EFAULT]	<i>abs_timeout</i> references invalid memory.
[EIDRM]	Specified message queue was removed during required operation.
[EINTR]	The <b>mq_receive()</b> or <b>mq_timedreceive()</b> operation was interrupted by a signal.
[EINVAL]	The process or thread would have blocked, and the <i>abs_timeout</i> parameter specified a nanoseconds field value less than 0 or greater than or equal to 1000 million.
[EINVAL]	<i>msg_ptr</i> value was null.
[EMSGSIZE]	The specified message buffer size, <i>msg_len</i> , is less than the message size attribute of the message queue.
[ENOTSUP]	Function is not supported with checkpoint-restart'ed processes.
[ETIMEDOUT]	The O_NONBLOCK flag was not set when the message queue was opened, but no message arrived on the queue before the specified timeout expired.

The `mq_receive()` and `mq_timedreceive()` functions might fail if:

[EBADMSG]                      The implementation has detected a data corruption problem with the message.

## Related Information

“mq\_send, mq\_timedsend Subroutine,” “msgctl Subroutine” on page 879, “msgget Subroutine” on page 881, “msgrcv Subroutine” on page 882, “msgsnd Subroutine” on page 885, “posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event Subroutine” on page 1163, “pthread\_mutex\_timedlock Subroutine” on page 1272, “pthread\_rwlock\_timedrdlock Subroutine” on page 1287, “pthread\_rwlock\_timedwrlock Subroutine” on page 1289.

The `sem_timedwait` and `time` subroutines in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The `mqueue.h` and `time.h` file.

---

## mq\_send, mq\_timedsend Subroutine

### Purpose

Sends a message to a message queue (REALTIME).

### Syntax

```
#include <mqueue.h>
```

```
int mq_send(mqd_t mqdes, const char *msg_ptr,  
            size_t msg_len, unsigned *msg_prio,
```

```
#include <mqueue.h>  
#include <time.h>
```

```
int mq_timedsend(mqd_t mqdes, const char *msg_ptr,  
                 size_t msg_len, unsigned msg_prio,  
                 const struct timespec *abs_timeout);
```

### Description

The `mq_send()` function adds the message pointed to by the argument `msg_ptr` to the message queue specified by `mqdes`. The `msg_len` argument specifies the length of the message, in bytes, pointed to by `msg_ptr`. The value of `msg_len` is less than or equal to the `mq_msgsize` attribute of the message queue, or `mq_send()` fails.

If the specified message queue is not full, `mq_send()` behaves as if the message is inserted into the message queue at the position indicated by the `msg_prio` argument. A message with a larger numeric value of `msg_prio` is inserted before messages with lower values of `msg_prio`. A message is inserted after other messages in the queue, if any, with equal `msg_prio` values. The value of `msg_prio` is less than {MQ\_PRIO\_MAX}.

If the specified message queue is full and `O_NONBLOCK` is not set in the message queue description associated with `mqdes`, `mq_send()` blocks until space becomes available to enqueue the message, or until `mq_send()` is interrupted by a signal. If more than one thread is waiting to send when space becomes available in the message queue and the **Priority Scheduling** option is supported, then the thread of the highest priority that has been waiting the longest is unblocked to send its message. Otherwise, it is unspecified which waiting thread is unblocked. If the specified message queue is full and `O_NONBLOCK` is set in the message queue description associated with `mqdes`, the message is not queued and `mq_send()` returns an error.



The **mq\_timedsend()** function adds a message to the message queue specified by *mqdes* in the manner defined for the **mq\_send()** function. However, if the specified message queue is full and `O_NONBLOCK` is not set in the message queue description associated with *mqdes*, the wait for sufficient room in the queue is terminated when the specified timeout expires. If `O_NONBLOCK` is set in the message queue description, this function matches **mq\_send()**.

The timeout expires when the absolute time specified by *abs\_timeout* passes—as measured by the clock on which timeouts are based (that is, when the value of that clock equals or exceeds *abs\_timeout*)—or when the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

If the **Timers** option is supported, the timeout is based on the `CLOCK_REALTIME` clock; if the **Timers** option is not supported, the timeout is based on the system clock as returned by the **time()** function.

The operation never fails with a timeout if there is sufficient room in the queue to add the message immediately. The validity of the *abs\_timeout* parameter does not need to be checked when there is sufficient room in the queue.

## Application Usage

The value of the symbol `{MQ_PRIO_MAX}` limits the number of priority levels supported by the application. Message priorities range from 0 to `{MQ_PRIO_MAX}-1`.

## Return Values

Upon successful completion, the **mq\_send()** and **mq\_timedsend()** functions return a value of 0. Otherwise, no message is enqueued, the functions return -1, and *errno* is set to indicate the error.

## Error Codes

The **mq\_send()** and **mq\_timedsend()** functions fail if:

[EAGAIN]	The <code>O_NONBLOCK</code> flag is set in the message queue description associated with <i>mqdes</i> , and the specified message queue is full.
[EBADF]	The <i>mqdes</i> argument is not a valid message queue descriptor open for writing.
[EFAULT]	<i>abs_timeout</i> references invalid memory.
[EIDRM]	Specified message queue was removed during required operation.
[EINTR]	A signal interrupted the call to <b>mq_send()</b> or <b>mq_timedsend()</b> .
[EINVAL]	The value of <i>msg_prio</i> was outside the valid range.
[EINVAL]	<i>msg_ptr</i> value was null.
[EINVAL]	The process or thread would have blocked, and the <i>abs_timeout</i> parameter specified a nanoseconds field value less than 0 or greater than or equal to 1000 million.
[EMSGSIZE]	The specified message length, <i>msg_len</i> , exceeds the message size attribute of the message queue.
[ENOTSUP]	Function is not supported with checkpoint-restart'ed processes.
[ETIMEDOUT]	The <code>O_NONBLOCK</code> flag was not set when the message queue was opened, but the timeout expired before the message could be added to the queue.

The **mq\_send()** and **mq\_timedsend()** functions might fail if:

[EBADMSG]	The implementation has detected a data corruption problem with the message.
-----------	---

## Related Information

“mq\_receive, mq\_timedreceive Subroutine” on page 870, “msgctl Subroutine” on page 879, “msgget Subroutine” on page 881, “msgrcv Subroutine” on page 882, “msgsnd Subroutine” on page 885, “posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event Subroutine” on page 1163

page 1163, “pthread\_mutex\_timedlock Subroutine” on page 1272, “pthread\_rwlock\_timedrdlock Subroutine” on page 1287, “pthread\_rwlock\_timedwrlock Subroutine” on page 1289.

The **sem\_timedwait** and **time** subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The **mqueue.h** and **time.h** file.

---

## mq\_unlink Subroutine

### Purpose

Removes a message queue.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mqueue.h>

int mq_unlink (name)
const char *name;
```

### Description

The **mq\_unlink** subroutine removes the message queue named by the pathname *name*. After a successful call to the **mq\_unlink** subroutine with the *name* parameter, a call to the **mq\_open** subroutine with the *name* parameter and the **O\_CREAT** flag will create a new message queue. If one or more processes have the message queue open when the **mq\_unlink** subroutine is called, destruction of the message queue is postponed until all references to the message queue have been closed.

After a successful completion of the **mq\_unlink** subroutine, calls to the **mq\_open** subroutine to recreate a message queue with the same name will succeed. The **mq\_unlink** subroutine never blocks even if all references to the message queue have not been closed.

### Parameters

*name* Specifies the message queue to be removed.

### Return Values

Upon successful completion, the **mq\_unlink** subroutine returns a zero. Otherwise, the named message queue is unchanged, and the **mq\_unlink** subroutine returns a -1 and sets **errno** to indicate the error.

### Error Codes

The **mq\_unlink** subroutine fails if:

<b>EACCES</b>	Permission is denied to unlink the named message queue.
<b>EFAULT</b>	Invalid used address.
<b>EINVAL</b>	The <i>name</i> parameter value is not valid
<b>ENAMETOOLONG</b>	The length of the <i>name</i> parameter exceeds <b>PATH_MAX</b> or a pathname component is longer than <b>NAME_MAX</b> .
<b>ENOENT</b>	The named message queue does not exist.
<b>ENOTSUP</b>	This function is not supported with processes that have been checkpoint-restart'ed.

## Related Information

“mq\_open Subroutine” on page 864 and “mq\_close Subroutine” on page 861.

---

## msem\_init Subroutine

### Purpose

Initializes a semaphore in a mapped file or shared memory region.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
msemaphore *msem_init ( Sem, InitialValue)  
msemaphore *Sem;  
int InitialValue;
```

### Description

The **msem\_init** subroutine allocates a new binary semaphore and initializes the state of the new semaphore.

If the value of the *InitialValue* parameter is **MSEM\_LOCKED**, the new semaphore is initialized in the locked state. If the value of the *InitialValue* parameter is **MSEM\_UNLOCKED**, the new semaphore is initialized in the unlocked state.

The **msemaphore** structure is located within a mapped file or shared memory region created by a successful call to the **mmap** subroutine and having both read and write access.

Whether a semaphore is created in a mapped file or in an anonymous shared memory region, any reference by a process that has mapped the same file or shared region, using an **msemaphore** structure pointer that resolved to the same file or start of region offset, is taken as a reference to the same semaphore.

Any previous semaphore state stored in the **msemaphore** structure is ignored and overwritten.

### Parameters

<i>Sem</i>	Points to an <b>msemaphore</b> structure in which the state of the semaphore is stored.
<i>Initial Value</i>	Determines whether the semaphore is locked or unlocked at allocation.

### Return Values

When successful, the **msem\_init** subroutine returns a pointer to the initialized **msemaphore** structure. Otherwise, it returns a null value and sets the **errno** global variable to indicate the error.

### Error Codes

If the **msem\_init** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

<b>EINVAL</b>	Indicates the <i>InitialValue</i> parameter is not valid.
<b>ENOMEM</b>	Indicates a new semaphore could not be created.

## Related Information

The **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine, **msem\_lock** (“msem\_lock Subroutine”) subroutine, **msem\_remove** (“msem\_remove Subroutine” on page 877) subroutine, **msem\_unlock** (“msem\_unlock Subroutine” on page 878) subroutine.

List of Memory Mapping Services and Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## msem\_lock Subroutine

### Purpose

Locks a semaphore.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
int msem_lock ( Sem, Condition)
msemaphore *Sem;
int Condition;
```

### Description

The **msem\_lock** subroutine attempts to lock a binary semaphore.

If the semaphore is not currently locked, it is locked and the **msem\_lock** subroutine completes successfully.

If the semaphore is currently locked, and the value of the *Condition* parameter is **MSEM\_IF\_NOWAIT**, the **msem\_lock** subroutine returns with an error. If the semaphore is currently locked, and the value of the *Condition* parameter is 0, the **msem\_lock** subroutine does not return until either the calling process is able to successfully lock the semaphore or an error condition occurs.

All calls to the **msem\_lock** and **msem\_unlock** subroutines by multiple processes sharing a common **msemaphore** structure behave as if the call were serialized.

If the **msemaphore** structure contains any value not resulting from a call to the **msem\_init** subroutine, followed by a (possibly empty) sequence of calls to the **msem\_lock** and **msem\_unlock** subroutines, the results are undefined. The address of an **msemaphore** structure is significant. If the **msemaphore** structure contains any value copied from an **msemaphore** structure at a different address, the result is undefined.

### Parameters

<i>Sem</i>	Points to an <b>msemaphore</b> structure that specifies the semaphore to be locked.
<i>Condition</i>	Determines whether the <b>msem_lock</b> subroutine waits for a currently locked semaphore to unlock.

### Return Values

When successful, the **msem\_lock** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **msem\_lock** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

- EAGAIN** Indicates a value of **MSEM\_IF\_NOWAIT** is specified for the *Condition* parameter and the semaphore is already locked.
- EINVAL** Indicates the *Sem* parameter points to an **msemaphore** structure specifying a semaphore that has been removed, or the *Condition* parameter is invalid.
- EINTR** Indicates the **msem\_lock** subroutine was interrupted by a signal that was caught.

## Related Information

The **msem\_init** (“msem\_init Subroutine” on page 875) subroutine, **msem\_remove** (“msem\_remove Subroutine”) subroutine, **msem\_unlock** (“msem\_unlock Subroutine” on page 878) subroutine.

List of Memory Mapping Services and Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## msem\_remove Subroutine

### Purpose

Removes a semaphore.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
int msem_remove ( Sem)  
msemaphore *Sem;
```

### Description

The **msem\_remove** subroutine removes a binary semaphore. Any subsequent use of the **msemaphore** structure before it is again initialized by calling the **msem\_init** subroutine will have undefined results.

The **msem\_remove** subroutine also causes any process waiting in the **msem\_lock** subroutine on the removed semaphore to return with an error.

If the **msemaphore** structure contains any value not resulting from a call to the **msem\_init** subroutine, followed by a (possibly empty) sequence of calls to the **msem\_lock** and **msem\_unlock** subroutines, the result is undefined. The address of an **msemaphore** structure is significant. If the **msemaphore** structure contains any value copied from an **msemaphore** structure at a different address, the result is undefined.

### Parameters

*Sem* Points to an **msemaphore** structure that specifies the semaphore to be removed.

### Return Values

When successful, the **msem\_remove** subroutine returns a value of 0. Otherwise, it returns a -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **msem\_remove** subroutine is unsuccessful, the **errno** global variable is set to the following value:

**EINVAL** Indicates the *Sem* parameter points to an **msemaphore** structure that specifies a semaphore that has been removed.

## Related Information

The **msem\_init** (“msem\_init Subroutine” on page 875) subroutine, **msem\_lock** (“msem\_lock Subroutine” on page 876) subroutine, **msem\_unlock** (“msem\_unlock Subroutine”) subroutine.

List of Memory Mapping Services and Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## msem\_unlock Subroutine

### Purpose

Unlocks a semaphore.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
int msem_unlock ( Sem, Condition)
msemaphore *Sem;
int Condition;
```

### Description

The **msem\_unlock** subroutine attempts to unlock a binary semaphore.

If the semaphore is currently locked, it is unlocked and the **msem\_unlock** subroutine completes successfully.

If the *Condition* parameter is 0, the semaphore is unlocked, regardless of whether or not any other processes are currently attempting to lock it. If the *Condition* parameter is set to the **MSEM\_IF\_WAITERS** value, and another process is waiting to lock the semaphore or it cannot be reliably determined whether some process is waiting to lock the semaphore, the semaphore is unlocked by the calling process. If the *Condition* parameter is set to the **MSEM\_IF\_WAITERS** value and no process is waiting to lock the semaphore, the semaphore will not be unlocked and an error will be returned.

### Parameters

<i>Sem</i>	Points to an <b>msemaphore</b> structure that specifies the semaphore to be unlocked.
<i>Condition</i>	Determines whether the <b>msem_unlock</b> subroutine unlocks the semaphore if no other processes are waiting to lock it.

### Return Values

When successful, the **msem\_unlock** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **msem\_unlock** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

- EAGAIN** Indicates a *Condition* value of **MSEM\_IF\_WAITERS** was specified and there were no waiters.  
**EINVAL** Indicates the *Sem* parameter points to an **msemaphore** structure specifying a semaphore that has been removed, or the *Condition* parameter is not valid.

## Related Information

The **msem\_init** (“msem\_init Subroutine” on page 875) subroutine, **msem\_lock** (“msem\_lock Subroutine” on page 876) subroutine, **msem\_remove** (“msem\_remove Subroutine” on page 877) subroutine.

List of Memory Mapping Services and Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## msgctl Subroutine

### Purpose

Provides message control operations.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/msg.h>
```

```
int msgctl (MessageQueueID, Command, Buffer)
int MessageQueueID, Command;
struct msqid_ds * Buffer;
```

### Description

The **msgctl** subroutine provides a variety of message control operations as specified by the *Command* parameter and stored in the structure pointed to by the *Buffer* parameter. The **msqid\_ds** structure is defined in the **sys/msg.h** file.

The following limits apply to the message queue:

- Maximum message size is 65,535 bytes for releases prior to AIX 4.1.5 and is 4 Megabytes for release AIX 4.1.5 and later releases.
- Maximum number of messages per queue is 524288.
- Maximum number of message queue IDs is 4096 for releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.
- Maximum number of bytes in a queue is 4 65,535 for releases prior to AIX 4.1.5 and is 4 Megabytes for release 4.1.5 and later releases.

### Parameters

*MessageQueueID* Specifies the message queue identifier.

*Command*

The following values for the *Command* parameter are available:

#### **IPC\_STAT**

Stores the current value of the above fields of the data structure associated with the *MessageQueueID* parameter into the **msqid\_ds** structure pointed to by the *Buffer* parameter.

The current process must have read permission in order to perform this operation.

#### **IPC\_SET**

Sets the value of the following fields of the data structure associated with the *MessageQueueID* parameter to the corresponding values found in the structure pointed to by the *Buffer* parameter:

```
msg_perm.uid  
msg_perm.gid  
msg_perm.mode/*Only the low-order  
nine bits*/  
msg_qbytes
```

The effective user ID of the current process must have root user authority or must equal the value of the `msg_perm.uid` or `msg_perm.cuid` field in the data structure associated with the *MessageQueueID* parameter in order to perform this operation. To raise the value of the `msg_qbytes` field, the effective user ID of the current process must have root user authority.

#### **IPC\_RMID**

Removes the message queue identifier specified by the *MessageQueueID* parameter from the system and destroys the message queue and data structure associated with it. The effective user ID of the current process must have root user authority or be equal to the value of the `msg_perm.uid` or `msg_perm.cuid` field in the data structure associated with the *MessageQueueID* parameter to perform this operation.

*Buffer*

Points to a **msqid\_ds** structure.

## Return Values

Upon successful completion, the **msgctl** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **msgctl** subroutine is unsuccessful if any of the following conditions is true:

- EINVAL** The *Command* or *MessageQueueID* parameter is not valid.
- EACCES** The *Command* parameter is equal to the **IPC\_STAT** value, and the calling process was denied read permission.
- EPERM** The *Command* parameter is equal to the **IPC\_RMID** value and the effective user ID of the calling process does not have root user authority. Or, the *Command* parameter is equal to the **IPC\_SET** value, and the effective user ID of the calling process is not equal to the value of the `msg_perm.uid` field or the `msg_perm.cuid` field in the data structure associated with the *MessageQueueID* parameter.
- EPERM** The *Command* parameter is equal to the **IPC\_SET** value, an attempt was made to increase the value of the `msg_qbytes` field, and the effective user ID of the calling process does not have root user authority.
- EFAULT** The *Buffer* parameter points outside of the process address space.

## Related Information

The **msgget** (“msgget Subroutine” on page 881) subroutine, **msgrcv** (“msgrcv Subroutine” on page 882) subroutine, **msgsnd** (“msgsnd Subroutine” on page 885) subroutine, **msgxrcv** (“msgxrcv Subroutine” on page 887) subroutine.



---

## msgget Subroutine

### Purpose

Gets a message queue identifier.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/msg.h>
```

```
int msgget ( Key, MessageFlag)
```

```
key_t Key;
```

```
int MessageFlag;
```

### Description

The **msgget** subroutine returns the message queue identifier associated with the specified *Key* parameter.

A message queue identifier, associated message queue, and data structure are created for the value of the *Key* parameter if one of the following conditions is true:

- The *Key* parameter is equal to the **IPC\_PRIVATE** value.
- The *Key* parameter does not already have a message queue identifier associated with it, and the **IPC\_CREAT** value is set.

Upon creation, the data structure associated with the new message queue identifier is initialized as follows:

- The *msg\_perm.cuid*, *msg\_perm.uid*, *msg\_perm.cgid*, and *msg\_perm.gid* fields are set equal to the effective user ID and effective group ID, respectively, of the calling process.
- The low-order 9 bits of the *msg\_perm.mode* field are set equal to the low-order 9 bits of the *MessageFlag* parameter.
- The *msg\_qnum*, *msg\_lspid*, *msg\_lrpid*, *msg\_stime*, and *msg\_rtime* fields are set equal to 0.
- The *msg\_ctime* field is set equal to the current time.
- The *msg\_qbytes* field is set equal to the system limit.

The **msgget** subroutine performs the following actions:

- The **msgget** subroutine either finds or creates (depending on the value of the *MessageFlag* parameter) a queue with the *Key* parameter.
- The **msgget** subroutine returns the ID of the queue header to its caller.

Limits on message size and number of messages in the queue can be found in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

### Parameters

*Key* Specifies either the value **IPC\_PRIVATE** or an Interprocess Communication (IPC) key constructed by the **ftok** (“ftok Subroutine” on page 322) subroutine (or by a similar algorithm).

*MessageFlag*

Constructed by logically ORing one or more of the following values:

**IPC\_CREAT**

Creates the data structure if it does not already exist.

**IPC\_EXCL**

Causes the **msgget** subroutine to fail if the **IPC\_CREAT** value is also set and the data structure already exists.

**S\_IRUSR**

Permits the process that owns the data structure to read it.

**S\_IWUSR**

Permits the process that owns the data structure to modify it.

**S\_IRGRP**

Permits the group associated with the data structure to read it.

**S\_IWGRP**

Permits the group associated with the data structure to modify it.

**S\_IROTH**

Permits others to read the data structure.

**S\_IWOTH**

Permits others to modify the data structure.

Values that begin with **S\_I** are defined in the **sys/mode.h** file and are a subset of the access permissions that apply to files.

## Return Values

Upon successful completion, the **msgget** subroutine returns a message queue identifier. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **msgget** subroutine is unsuccessful if any of the following conditions is true:

- EACCES** A message queue identifier exists for the *Key* parameter, but operation permission as specified by the low-order 9 bits of the *MessageFlag* parameter is not granted.
- ENOENT** A message queue identifier does not exist for the *Key* parameter and the **IPC\_CREAT** value is not set.
- ENOSPC** A message queue identifier is to be created, but the system-imposed limit on the maximum number of allowed message queue identifiers system-wide would be exceeded.
- EEXIST** A message queue identifier exists for the *Key* parameter, and both **IPC\_CREAT** and **IPC\_EXCL** values are set.

## Related Information

The **ftok** (“ftok Subroutine” on page 322) subroutine, **msgctl** (“msgctl Subroutine” on page 879) subroutine, **msgrcv** (“msgrcv Subroutine”) subroutine, **msgsnd** (“msgsnd Subroutine” on page 885) subroutine, **msgxrcv** (“msgxrcv Subroutine” on page 887) subroutine.

The **mode.h** file.

---

## msgrcv Subroutine

### Purpose

Reads a message from a queue.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/msg.h>
```

```
int msgrcv (MessageQueueID, MessagePointer, MessageSize, MessageType, MessageFlag)
int MessageQueueID, MessageFlag;
void * MessagePointer;
size_t MessageSize;
long int MessageType;
```

## Description

The **msgrcv** subroutine reads a message from the queue specified by the *MessageQueueID* parameter and stores it into the structure pointed to by the *MessagePointer* parameter. The current process must have read permission in order to perform this operation.

**Note:** The routine may coredump instead of returning EFAULT when an invalid pointer is passed in case of 64-bit application calling 32-bit kernel interface.

Limits on message size and number of messages in the queue can be found in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

**Note:** For a 64-bit process, the **mtype** field is 64 bits long. However, for compatibility with 32-bit processes, the **mtype** field must be a 32-bit signed value that is sign-extended to 64 bits. The most significant 32 bits are not put on the message queue. For a 64-bit process, the **mtype** field is again sign-extended to 64 bits.

## Parameters

<i>MessageQueueID</i>	Specifies the message queue identifier.
<i>MessagePointer</i>	Points to a <b>msgbuf</b> structure containing the message. The <b>msgbuf</b> structure is defined in the <b>sys/msg.h</b> file and contains the following fields: <pre>mtyp_t  mtype;          /* Message type */ char    mtext[1];      /* Beginning of message text */</pre> <p>The <b>mtype</b> field contains the type of the received message as specified by the sending process. The <b>mtext</b> field is the text of the message.</p>
<i>MessageSize</i>	Specifies the size of the <b>mtext</b> field in bytes. The received message is truncated to the size specified by the <i>MessageSize</i> parameter if it is longer than the size specified by the <i>MessageSize</i> parameter and if the <b>MSG_NOERROR</b> value is set in the <i>MessageFlag</i> parameter. The truncated part of the message is lost and no indication of the truncation is given to the calling process.
<i>MessageType</i>	Specifies the type of message requested as follows: <ul style="list-style-type: none"><li>• If equal to the value of 0, the first message on the queue is received.</li><li>• If greater than 0, the first message of the type specified by the <i>MessageType</i> parameter is received.</li><li>• If less than 0, the first message of the lowest type that is less than or equal to the absolute value of the <i>MessageType</i> parameter is received.</li></ul>

*MessageFlag* Specifies either a value of 0 or is constructed by logically ORing one or more of the following values:

**MSG\_NOERROR**

Truncates the message if it is longer than the *MessageSize* parameter.

**IPC\_NOWAIT**

Specifies the action to take if a message of the desired type is not on the queue:

- If the **IPC\_NOWAIT** value is set, the calling process returns a value of -1 and sets the **errno** global variable to the **ENOMSG** error code.
- If the **IPC\_NOWAIT** value is not set, the calling process suspends execution until one of the following occurs:
  - A message of the desired type is placed on the queue.
  - The message queue identifier specified by the *MessageQueueID* parameter is removed from the system. When this occurs, the **errno** global variable is set to the **EIDRM** error code, and a value of -1 is returned.
  - The calling process receives a signal that is to be caught. In this case, a message is not received and the calling process resumes in the manner described in the **sigaction** subroutine.

## Return Values

Upon successful completion, the **msgrcv** subroutine returns a value equal to the number of bytes actually stored into the *mtext* field and the following actions are taken with respect to fields of the data structure associated with the *MessageQueueID* parameter:

- The *msg\_qnum* field is decremented by 1.
- The *msg\_lrpid* field is set equal to the process ID of the calling process.
- The *msg\_rtime* field is set equal to the current time.

If the **msgrcv** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **msgrcv** subroutine is unsuccessful if any of the following conditions is true:

<b>EINVAL</b>	The <i>MessageQueueID</i> parameter is not a valid message queue identifier.
<b>EACCES</b>	The calling process is denied permission for the specified operation.
<b>E2BIG</b>	The <i>mtext</i> field is greater than the <i>MessageSize</i> parameter, and the <b>MSG_NOERROR</b> value is not set.
<b>ENOMSG</b>	The queue does not contain a message of the desired type and the <b>IPC_NOWAIT</b> value is set.
<b>EFAULT</b>	The <i>MessagePointer</i> parameter points outside of the allocated address space of the process.
<b>EINTR</b>	The <b>msgrcv</b> subroutine is interrupted by a signal.
<b>EIDRM</b>	The message queue identifier specified by the <i>MessageQueueID</i> parameter has been removed from the system.

## Related Information

The **msgctl** (“msgctl Subroutine” on page 879) subroutine, **msgget** (“msgget Subroutine” on page 881) subroutine, **msgsnd** (“msgsnd Subroutine” on page 885) subroutine, **msgxrcv** (“msgxrcv Subroutine” on page 887) subroutine, **sigaction** subroutine.

---

## msgsnd Subroutine

### Purpose

Sends a message.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/msg.h>
```

```
int msgsnd (MessageQueueID, MessagePointer, MessageSize, MessageFlag)
int MessageQueueID, MessageFlag;
const void * MessagePointer;
size_t MessageSize;
```

### Description

The **msgsnd** subroutine sends a message to the queue specified by the *MessageQueueID* parameter. The current process must have write permission to perform this operation. The *MessagePointer* parameter points to an **msgbuf** structure containing the message. The **sys/msg.h** file defines the **msgbuf** structure. The structure contains the following fields:

```
mtyp_t mtype; /* Message type */
char mtext[1]; /* Beginning of message text */
```

The *mtype* field specifies a positive integer used by the receiving process for message selection. The *mtext* field can be any text of the length in bytes specified by the *MessageSize* parameter. The *MessageSize* parameter can range from 0 to the maximum limit imposed by the system.

The following example shows a typical user-defined **msgbuf** structure that includes sufficient space for the largest message:

```
struct my_msgbuf
mtyp_t mtype;
char mtext[MSGSZ]; /* MSGSZ is the size of the largest message */
```

**Note:** The routine may coredump instead of returning EFAULT when an invalid pointer is passed in case of 64-bit application calling 32-bit kernel interface.

The following system limits apply to the message queue:

- Maximum message size is 65,535 bytes for releases prior to AIX 4.1.5 and is 4 Megabytes for AIX 4.1.5 and later releases.
- Maximum number of messages per queue is 524288.
- Maximum number of message queue IDs is 4096 for releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.
- Maximum number of bytes in a queue is 4 65,535 bytes for releases prior to AIX 4.1.5 is 4 Megabytes for AIX 4.1.5 and later releases.

**Note:** For a 64-bit process, the **mtype** field is 64 bits long. However, for compatibility with 32-bit processes, the **mtype** field must be a 32-bit signed value that is sign-extended to 64 bits. The most significant 32 bits are not put on the message queue. For a 64-bit process, the **mtype** field is again sign-extended to 64 bits.

The *MessageFlag* parameter specifies the action to be taken if the message cannot be sent for one of the following reasons:

- The number of bytes already on the queue is equal to the number of bytes defined by the **msg\_qbytes** structure.
- The total number of messages on the queue is equal to a system-imposed limit.

These actions are as follows:

- If the *MessageFlag* parameter is set to the **IPC\_NOWAIT** value, the message is not sent, and the **msgsnd** subroutine returns a value of -1 and sets the **errno** global variable to the **EAGAIN** error code.
- If the *MessageFlag* parameter is set to 0, the calling process suspends execution until one of the following occurs:
  - The condition responsible for the suspension no longer exists, in which case the message is sent.
  - The *MessageQueueID* parameter is removed from the system. (For information on how to remove the *MessageQueueID* parameter, see the **msgctl** (“msgctl Subroutine” on page 879) subroutine.) When this occurs, the **errno** global variable is set equal to the **EIDRM** error code, and a value of -1 is returned.
  - The calling process receives a signal that is to be caught. In this case the message is not sent and the calling process resumes execution in the manner prescribed in the **sigaction** subroutine.

## Parameters

<i>MessageQueueID</i>	Specifies the queue to which the message is sent.
<i>MessagePointer</i>	Points to a <b>msgbuf</b> structure containing the message.
<i>MessageSize</i>	Specifies the length, in bytes, of the message text.
<i>MessageFlag</i>	Specifies the action to be taken if the message cannot be sent.

## Return Values

Upon successful completion, a value of 0 is returned and the following actions are taken with respect to the data structure associated with the *MessageQueueID* parameter:

- The **msg\_qnum** field is incremented by 1.
- The **msg\_lspid** field is set equal to the process ID of the calling process.
- The **msg\_stime** field is set equal to the current time.

If the **msgsnd** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **msgsnd** subroutine is unsuccessful and no message is sent if one or more of the following conditions is true:

<b>EACCES</b>	The calling process is denied permission for the specified operation.
<b>EAGAIN</b>	The message cannot be sent for one of the reasons stated previously, and the <i>MessageFlag</i> parameter is set to the <b>IPC_NOWAIT</b> value or the system has temporarily ran out of memory resource.
<b>EFAULT</b>	The <i>MessagePointer</i> parameter points outside of the address space of the process.
<b>EIDRM</b>	The message queue identifier specified by the <i>MessageQueueID</i> parameter has been removed from the system.
<b>EINTR</b>	The <b>msgsnd</b> subroutine received a signal.
<b>EINVAL</b>	The <i>MessageQueueID</i> parameter is not a valid message queue identifier.
<b>EINVAL</b>	The <b>mtype</b> field is less than 1.
<b>EINVAL</b>	The <i>MessageSize</i> parameter is less than 0 or greater than the system-imposed limit.
<b>EINVAL</b>	The upper 32-bits of the 64-bit <b>mtype</b> field for a 64-bit process is not 0.
<b>ENOMEM</b>	The message could not be sent because not enough storage space was available.

## Related Information

The **msgctl** (“msgctl Subroutine” on page 879) subroutine, **msgget** (“msgget Subroutine” on page 881) subroutine, **msgrcv** (“msgrcv Subroutine” on page 882) subroutine, **msgxrcv** (“msgxrcv Subroutine”) subroutine, **sigaction** subroutine.

---

## msgxrcv Subroutine

### Purpose

Receives an extended message.

### Library

Standard C Library (**libc.a**)

### Syntax

For releases prior to AIX 4.3:

```
#include <sys/msg.h>
```

```
int msgxrcv (MessageQueueID, MessagePointer, MessageSize, MessageType, MessageFlag)
int MessageQueueID, MessageFlag, MessageSize;
struct msgxbuf * MessagePointer;
long MessageType;
```

For AIX 4.3 and later releases:

```
#include <sys/msg.h>
```

```
int msgxrcv (MessageQueueID, MessagePointer, MessageSize, MessageType, MessageFlag)
int MessageQueueID, MessageFlag;
size_t MessageSize;
struct msgxbuf * MessagePointer;
long MessageType;
```

### Description

The **msgxrcv** subroutine reads a message from the queue specified by the *MessageQueueID* parameter and stores it into the extended message receive buffer pointed to by the *MessagePointer* parameter. The current process must have read permission in order to perform this operation. The **msgxbuf** structure is defined in the **sys/msg.h** file.

**Note:** The routine may coredump instead of returning EFAULT when an invalid pointer is passed in case of 64-bit application calling 32-bit kernel interface.

The following limits apply to the message queue:

- Maximum message size is 65,535 bytes for releases prior to AIX 4.1.5 and is 4 Megabytes for AIX 4.1.5 and later releases.
- Maximum number of messages per queue is 8192.
- Maximum number of message queue IDs is 4096 for releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.
- Maximum number of bytes in a queue is 4 65,535 for releases prior to AIX 4.1.5 and is 4 Megabytes for AIX 4.1.5 later releases.

**Note:** For a 64-bit process, the **mtype** field is 64 bits long. However, for compatibility with 32-bit processes, the **mtype** field must be a 32-bit signed value that is sign-extended to 64 bits. The most significant 32 bits are not put on the message queue. For a 64-bit process, the **mtype** field is again sign-extended to 64 bits.

## Parameters

<i>MessageQueueID</i>	Specifies the message queue identifier.
<i>MessagePointer</i>	Specifies a pointer to an extended message receive buffer where a message is stored.
<i>MessageSize</i>	Specifies the size of the <i>mtext</i> field in bytes. The receive message is truncated to the size specified by the <i>MessageSize</i> parameter if it is larger than the <i>MessageSize</i> parameter and the <b>MSG_NOERROR</b> value is true. The truncated part of the message is lost and no indication of the truncation is given to the calling process. If the message is longer than the number of bytes specified by the <i>MessageSize</i> parameter and the <b>MSG_NOERROR</b> value is not set, the <b>msgxrcv</b> subroutine is unsuccessful and sets the <b>errno</b> global variable to the <b>E2BIG</b> error code.
<i>MessageType</i>	Specifies the type of message requested as follows: <ul style="list-style-type: none"><li>• If the <i>MessageType</i> parameter is equal to 0, the first message on the queue is received.</li><li>• If the <i>MessageType</i> parameter is greater than 0, the first message of the type specified by the <i>MessageType</i> parameter is received.</li><li>• If the <i>MessageType</i> parameter is less than 0, the first message of the lowest type that is less than or equal to the absolute value of the <i>MessageType</i> parameter is received.</li></ul>
<i>MessageFlag</i>	Specifies a value of 0 or a value constructed by logically ORing one or more of the following values: <b>MSG_NOERROR</b> Truncates the message if it is longer than the number of bytes specified by the <i>MessageSize</i> parameter. <b>IPC_NOWAIT</b> Specifies the action to take if a message of the desired type is not on the queue: <ul style="list-style-type: none"><li>• If the <b>IPC_NOWAIT</b> value is set, the calling process returns a value of -1 and sets the <b>errno</b> global variable to the <b>ENOMSG</b> error code.</li><li>• If the <b>IPC_NOWAIT</b> value is not set, the calling process suspends execution until one of the following occurs:<ul style="list-style-type: none"><li>– A message of the desired type is placed on the queue.</li><li>– The message queue identifier specified by the <i>MessageQueueID</i> parameter is removed from the system. When this occurs, the <b>errno</b> global variable is set to the <b>EIDRM</b> error code, and a value of -1 is returned.</li><li>– The calling process receives a signal that is to be caught. In this case, a message is not received and the calling process resumes in the manner prescribed in the <b>sigaction</b> subroutine.</li></ul></li></ul>

## Return Values

Upon successful completion, the **msgxrcv** subroutine returns a value equal to the number of bytes actually stored into the *mtext* field, and the following actions are taken with respect to the data structure associated with the *MessageQueueID* parameter:

- The *msg\_qnum* field is decremented by 1.
- The *msg\_lrpId* field is set equal to the process ID of the calling process.
- The *msg\_rtime* field is set equal to the current time.

If the **msgxrcv** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.



## Error Codes

The **msgxrcv** subroutine is unsuccessful if any of the following conditions is true:

<b>EINVAL</b>	The <i>MessageQueueID</i> parameter is not a valid message queue identifier.
<b>EACCES</b>	The calling process is denied permission for the specified operation.
<b>EINVAL</b>	The <i>MessageSize</i> parameter is less than 0.
<b>E2BIG</b>	The <i>mtext</i> field is greater than the <i>MessageSize</i> parameter, and the <b>MSG_NOERROR</b> value is not set.
<b>ENOMSG</b>	The queue does not contain a message of the desired type and the <b>IPC_NOWAIT</b> value is set.
<b>EFAULT</b>	The <i>MessagePointer</i> parameter points outside of the process address space.
<b>EINTR</b>	The <b>msgxrcv</b> subroutine was interrupted by a signal.
<b>EIDRM</b>	The message queue identifier specified by the <i>MessageQueueID</i> parameter is removed from the system.

## Related Information

The **msgctl** (“msgctl Subroutine” on page 879) subroutine, **msgget** (“msgget Subroutine” on page 881) subroutine, **msgrcv** (“msgrcv Subroutine” on page 882) subroutine, **msgsnd** (“msgsnd Subroutine” on page 885) subroutine, **sigaction** subroutine.

---

## msleep Subroutine

### Purpose

Puts a process to sleep when a semaphore is busy.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/mman.h>
```

```
int msleep (Sem)  
msemaphore * Sem;
```

### Description

The **msleep** subroutine puts a calling process to sleep when a semaphore is busy. The semaphore should be located in a shared memory region. Use the **mmap** subroutine to create the shared memory section.

All of the values in the **msemaphore** structure must result from a **msem\_init** subroutine call. This call may or may not be followed by a sequence of calls to the **msem\_lock** subroutine or the **msem\_unlock** subroutine. If the **msemaphore** structure value originates in another manner, the results of the **msleep** subroutine are undefined.

The address of the **msemaphore** structure is significant. You should be careful not to modify the structure’s address. If the structure contains values copied from a **msemaphore** structure at another address, the results of the **msleep** subroutine are undefined.

### Parameters

*Sem* Points to the **msemaphore** structure that specifies the semaphore.

## Error Codes

If the **msleep** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

- EFAULT** Indicates that the *Sem* parameter points to an invalid address or the address does not contain a valid **msemaphore** structure.
- EINTR** Indicates that the process calling the **msleep** subroutine was interrupted by a signal while sleeping.

## Related Information

The **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine, **msem\_init** (“msem\_init Subroutine” on page 875) subroutine, **msem\_lock** (“msem\_lock Subroutine” on page 876) subroutine, **msem\_unlock** (“msem\_unlock Subroutine” on page 878) subroutine, **mwakeup** (“mwakeup Subroutine” on page 894) subroutine.

Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## msync Subroutine

### Purpose

Synchronize memory with physical storage.

### Library

Standard C Library (**libc.a**).

### Syntax

```
#include <sys/types.h>
#include <sys/mman.h>
```

```
int msync ( addr, len, flags)
void *addr;
size_t len;
int flags;
```

### Description

The **msync** subroutine controls the caching operations of a mapped file or shared memory region. Use the **msync** subroutine to transfer modified pages in the region to the underlying file storage device.

If the application has requested Single UNIX Specification, Version 2 compliant behavior then the **st\_ctime** and **st\_mtime** fields of the mapped file are marked for update upon successful completion of the **msync** subroutine call if the file has been modified.

### Parameters

- addr* Specifies the address of the region to be synchronized. Must be a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.
- len* Specifies the length, in bytes, of the region to be synchronized. If the *len* parameter is not a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter, the length of the region is rounded up to the next multiple of the page size.

*flags* Specifies one or more of the following symbolic constants that determine the way caching operations are performed:

#### **MS\_SYNC**

Specifies synchronous cache flush. The **msync** subroutine does not return until the system completes all I/O operations.

This flag is invalid when the **MAP\_PRIVATE** flag is used with the **mmap** subroutine.

**MAP\_PRIVATE** is the default privacy setting. When the **MS\_SYNC** and **MAP\_PRIVATE** flags both are used, the **msync** subroutine returns an **errno** value of **EINVAL**.

#### **MS\_ASYNC**

Specifies an asynchronous cache flush. The **msync** subroutine returns after the system schedules all I/O operations.

This flag is invalid when the **MAP\_PRIVATE** flag is used with the **mmap** subroutine.

**MAP\_PRIVATE** is the default privacy setting. When the **MS\_ASYNC** and **MAP\_PRIVATE** flags both are used, the **msync** subroutine returns an **errno** value of **EINVAL**.

#### **MS\_INVALIDATE**

Specifies that the **msync** subroutine invalidates all cached copies of the pages. New copies of the pages must then be obtained from the file system the next time they are referenced.

## Return Values

When successful, the **msync** subroutine returns 0. Otherwise, it returns -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **msync** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

<b>EBUSY</b>	One or more pages in the range passed to the <b>msync</b> subroutine is pinned.
<b>EIO</b>	An I/O error occurred while reading from or writing to the file system.
<b>ENOMEM</b>	The range specified by ( <i>addr</i> , <i>addr + len</i> ) is invalid for a process' address space, or the range specifies one or more unmapped pages.
<b>EINVAL</b>	The <i>addr</i> argument is not a multiple of the page size as returned by the <b>sysconf</b> subroutine using the <b>_SC_PAGE_SIZE</b> value for the <i>Name</i> parameter, or the <i>flags</i> parameter is invalid. The address of the region is within the process' inheritable address space.

---

## mt\_\_trce Subroutine

### Purpose

Dumps traceback information into a lightweight core file.

### Library

PTools Library (**libptools\_ptr.a**)

### Syntax

```
void mt__trce (int FileDescriptor, int Signal, struct sigcontext *Context, int Node);
```

### Description

The **mt\_\_trce** subroutine dumps traceback information of the calling thread and all other threads allocated in the process space into the file specified by the *FileDescriptor* parameter. The format of the output from this subroutine complies with the Parallel Tools Consortium Lightweight CoreFile Format. Threads, except the calling thread, will be suspended after the calling thread enters this subroutine and while the traceback information is being obtained. Threads execution resumes when this subroutine returns.

When using the **mt\_\_trce** subroutine in a signal handler, it is recommended that the application be started with the environment variable **AIXTHREAD\_SCOPE** set to **S** (As in `export AIXTHREAD_SCOPE=S`). If this variable is not set, the application may hang.

## Parameters

<i>Context</i>	Points to the <b>sigcontext</b> structure containing the context of the thread when the signal happens. The context is used to generate the traceback information for the calling thread. This is used only if the <i>Signal</i> parameter is nonzero. If the <b>mt__trce</b> subroutine is called with the <i>Signal</i> parameter set to zero, the <i>Context</i> parameter is ignored and the traceback information is generated based on the current context of the calling thread. Refer to the <b>sigaction</b> subroutine for further description about signal handlers and how the <b>sigcontext</b> structure is passed to a signal handler.
<i>File Descriptor</i>	The file descriptor of the lightweight core file. It specifies the target file into which the traceback information is written.
<i>Node</i>	Specifies the number of the tasks or nodes where this subroutine is executing and is used only for a parallel application consisting of multiple tasks. The <i>Node</i> parameter will be used in section headers of the traceback information to identify the task or node from which the information is generated.
<i>Signal</i>	The number of the signal that causes the signal handler to be executed. This is used only if the <b>mt__trce</b> subroutine is called from a signal handler. A Fault-Info section defined by the Parallel Tools Consortium Lightweight Core File Format will be written into the output lightweight core file based on this signal number. If the <b>mt__trce</b> subroutine is not called from a signal handler, the <i>Signal</i> parameter must be set to 0 and a Fault-Info section will not be generated.

### Notes:

1. To obtain source line information in the traceback, the programs must have been compiled with the **-g** option to include the necessary line number information in the executable files. Otherwise, address offset from the beginning of the function is provided.
2. Line number information is not provided for shared objects even if they were compiled with the **-g** option.
3. Function names are not provided if a program or a library is compiled with optimization. To obtain function name information in the traceback and still have the object code optimized, compiler option **-qtbtable=full** must be specified.
4. In rare cases, the traceback of a thread may seem to skip one level of procedure calls. This is because the traceback is obtained at the moment the thread entered a procedure and has not yet allocated a stack frame.

## Return Values

Upon successful completion, the **mt\_\_trce** subroutine returns a value of 0. Otherwise, an error number is returned to indicate the error.

## Error Codes

If an error occurs, the subroutine returns -1 and the **errno** global variable is set to indicate the error, as follows:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid file descriptor open for writing.
<b>ENOSPC</b>	No free space is left in the file system containing the file.
<b>EDQUOT</b>	New disk blocks cannot be allocated for the file because the user or group quota of blocks has been exhausted on the file system.
<b>EINVAL</b>	The value of the <i>Signal</i> parameter is invalid or the <i>Context</i> parameter points to an invalid context.
<b>ENOMEM</b>	Insufficient memory exists to perform the operation.

## Examples

1. The following example calls the `mt_trce` subroutine to generate traceback information in a signal handler.

```
void
my_handler(int signal,
           int code,
           struct sigcontext *sigcontext_data)
{
    int lcf_fd;
    ....
    lcf_fd = open(file_name, O_WRONLY|O_CREAT|O_APPEND, 0666);
    ....
    rc = mt_trce(lcf_fd, signal, sigcontext_data, 0);
    ....
    close(lcf_fd);
    ....
}
```

2. The following is an example of the lightweight core file generated by the `mt_trce` subroutine. Notice the thread ID in the information is the unique sequence number of a thread for the life time of the process containing the thread.

```
+++PARALLEL TOOLS CONSORTIUM LIGHTWEIGHT COREFILE FORMAT version 1.0
+++LCB 1.0 Thu Jun 30 16:02:35 1999 Generated by AIX
#
+++ID Node 0 Process 21084 Thread 1
***FAULT "SIGABRT - Abort"
+++STACK
func2 : 123 # in file
func1 : 272 # in file
main : 49 # in file
---STACK
---ID Node 0 Process 21084 Thread 1
#
+++ID Node 0 Process 21084 Thread 2
+++STACK
nsleep : 0x0000001c
sleep : 0x00000030
f_mt_exec : 21 # in file
_pthread_body : 0x00000114
---STACK
---ID Node 0 Process 21084 Thread 2
#
+++ID Node 0 Process 21084 Thread 3
+++STACK
nsleep : 0x0000001c
sleep : 0x00000030
f_mt_exec : 21 # in file
_pthread_body : 0x00000114
---STACK
---ID Node 0 Process 21084 Thread 3
---LCB
```

## Related Information

The `install_lwcf_handler` and `sigaction` subroutines.

---

## munmap Subroutine

### Purpose

Unmaps pages of memory.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/types.h>
#include <sys/mman.h>
```

```
int munmap ( addr, len)
void *addr;
size_t len;
```

## Description

The **munmap** subroutine unmaps a mapped file or shared memory region or anonymous memory region. The **munmap** subroutine unmaps regions created from calls to the **mmap** subroutine only.

If an address lies in a region that is unmapped by the **munmap** subroutine and that region is not subsequently mapped again, any reference to that address will result in the delivery of a **SIGSEGV** signal to the process.

## Parameters

*addr* Specifies the address of the region to be unmapped. Must be a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.

*len* Specifies the length, in bytes, of the region to be unmapped. If the *len* parameter is not a multiple of the page size returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter, the length of the region is rounded up to the next multiple of the page size.

## Return Values

When successful, the **munmap** subroutine returns 0. Otherwise, it returns -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **munmap** subroutine is unsuccessful, the **errno** global variable is set to the following value:

**EINVAL** The *addr* parameter is not a multiple of the page size as returned by the **sysconf** subroutine using the **\_SC\_PAGE\_SIZE** value for the *Name* parameter.

**EINVAL** The application has requested Single UNIX Specification, Version 2 compliant behavior and the *len* argument is 0.

---

## mwakeup Subroutine

### Purpose

Wakes up a process that is waiting on a semaphore.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/mman.h>
int mwaitup (Sem)
msemaphore * Sem;
```

## Description

The **mwakeup** subroutine wakes up a process that is sleeping and waiting for an idle semaphore. The semaphore should be located in a shared memory region. Use the **mmap** subroutine to create the shared memory section.

All of the values in the **msemaphore** structure must result from a **msem\_init** subroutine call. This call may or may not be followed by a sequence of calls to the **msem\_lock** subroutine or the **msem\_unlock** subroutine. If the **msemaphore** structure value originates in another manner, the results of the **mwakeup** subroutine are undefined.

The address of the **msemaphore** structure is significant. You should be careful not to modify the structure's address. If the structure contains values copied from a **msemaphore** structure at another address, the results of the **mwakeup** subroutine are undefined.

## Parameters

*Sem* Points to the **msemaphore** structure that specifies the semaphore.

## Return Values

When successful, the **mwakeup** subroutine returns a value of 0. Otherwise, this routine returns a value of -1 and sets the **errno** global variable to **EFAULT**.

## Error Codes

A value of **EFAULT** indicates that the *Sem* parameter points to an invalid address or that the address does not contain a valid **msemaphore** structure.

## Related Information

The **mmap** (“mmap or mmap64 Subroutine” on page 843) subroutine, **msem\_init** (“msem\_init Subroutine” on page 875) subroutine, **msem\_lock** (“msem\_lock Subroutine” on page 876) subroutine, **msem\_unlock** (“msem\_unlock Subroutine” on page 878) subroutine, and the **msleep** (“msleep Subroutine” on page 889) subroutine.

Understanding Memory Mapping in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## nan, nanf, nanl, nand32, nand64, and nand128 Subroutines

### Purpose

Return a quiet NaN.

### Syntax

```
#include <math.h>

double nan (tagp)
const char *tagp;

float nanf (tagp)
```

```
const char *tagp;

long double nanl (tagp)
const char *tagp;

_Decimal32 nand32(tagp)
const char *tagp;

_Decimal64 nand64(tagp)
const char *tagp;

_Decimal128 nand128(tagp)
const char *tagp;
```

## Description

The function call **nan**("*n-char-sequence*") is equivalent to:  
strtod("NAN(*n-char-sequence*)", (char \*\*) NULL);

The function call **nan**(" ") is equivalent to:  
strtod("NAN()", (char \*\*) NULL)

If *tagp* does not point to an *n-char* sequence or an empty string, the function call is equivalent to:  
strtod("NAN", (char \*\*) NULL)

Function calls to the **nanf**, **nanl**, **nand32**, **nand64**, and **nand128** subroutines are equivalent to the corresponding function calls to the **strtof**, **strtold**, **strtod32**, **strtod64**, and **strtod128** subroutines.

## Parameters

*tagp*            Indicates the content of the quiet NaN.

## Return Values

The **nan**, **nanf**, **nanl**, **nand32**, **nand64**, and **nand128** subroutines return a quiet NaN with content indicated through *tagp*.

## Related Information

The "atof atoff Subroutine" on page 96.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## nanosleep Subroutine

### Purpose

Causes the current thread to be suspended from execution.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <time.h>

int nanosleep (rqtp, rmtpt)
const struct timespec *rqtp;
struct timespec *rmtpt;
```



## Description

The **nanosleep** subroutine causes the current thread to be suspended from execution until either the time interval specified by the *rqtp* parameter has elapsed or a signal is delivered to the calling thread and its action is to invoke a signal-catching function or to terminate the process. The suspension time may be longer than requested because the argument value is rounded up to an integer multiple of the sleep resolution. This can also occur because of the scheduling of other activity by the system. Unless it is interrupted by a signal, the suspension time will not be less than the time specified by the *rqtp* parameter, as measured by the system clock **CLOCK\_REALTIME**.

The use of the **nanosleep** subroutine has no effect on the action or blockage of any signal.

## Parameters

*rqtp* Specifies the time interval that the thread is suspended.  
*rmtp* Points to the **timespec** structure.

## Return Values

If the **nanosleep** subroutine returns because the requested time has elapsed, its return value is zero.

If the **nanosleep** subroutine returns because it has been interrupted by a signal, it returns -1 and sets **errno** to indicate the interruption. If the *rmtp* parameter is non-NULL, the **timespec** structure is updated to contain the amount of time remaining in the interval (the requested time minus the time actually slept). If the *rmtp* parameter is NULL, the remaining time is not returned.

If the **nanosleep** subroutine fails, it returns -1 and sets **errno** to indicate the error.

## Error Codes

The **nanosleep** subroutine fails if:

**EINTR** The **nanosleep** subroutine was interrupted by a signal.  
**EINVAL** The *rqtp* parameter specified a nanosecond value less than zero or greater than or equal to 1000 million.

## Related Information

The sleep subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## nearbyint, nearbyintf, nearbyintl, nearbyintd32, nearbyintd64, and nearbyintd128 Subroutines

### Purpose

Round numbers to an integer value in floating-point format.

### Syntax

```
#include <math.h>
```

```
double nearbyint (x)  
double x;
```

```
float nearbyintf (x)  
float x;
```

**long double nearbyintl** (*x*)  
**long double** *x*;

**\_Decimal32 nearbyintd32**(*x*)  
**\_Decimal32** *x*;

**\_Decimal64 nearbyintd64**(*x*)  
**\_Decimal64** *x*;

**\_Decimal128 nearbyintd128**(*x*)  
**\_Decimal128** *x*;

## Description

The **nearbyint**, **nearbyintf**, **nearbyintl**, **nearbyintd32**, **nearbyintd64**, and **nearbyintd128** subroutines round the *x* parameter to an integer value in floating-point format, using the current rounding direction and without raising the inexact floating-point exception.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **nearbyint**, **nearbyintf**, **nearbyintl**, **nearbyintd32**, **nearbyintd64**, and **nearbyintd128** subroutines return the rounded integer value.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ ,  $\pm 0$  is returned.

If *x* is  $\pm \text{Inf}$ , *x* is returned.

If the correct value would cause overflow, a range error occurs and the **nearbyint**, **nearbyintf**, **nearbyintl**, **nearbyintd32**, **nearbyintd64**, and **nearbyintd128** subroutines return the value of the macro  **$\pm \text{HUGE\_VAL}$** ,  **$\pm \text{HUGE\_VALF}$** ,  **$\pm \text{HUGE\_VALL}$** ,  **$\pm \text{HUGE\_VAL\_D32}$** ,  **$\pm \text{HUGE\_VAL\_D64}$** ,  **$\pm \text{HUGE\_VAL\_D128}$**  (with the same sign as *x*), respectively.

## Related Information

“**feclearexcept** Subroutine” on page 265 and “**fetestexcept** Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## **nextafter**, **nextafterf**, **nextafterl**, **nexttoward**, **nexttowardf**, or **nexttowardl** Subroutine

### Purpose

Computes the next representable floating-point number.

## Syntax

```
#include <math.h>
```

```
float nextafterf (x, y)
```

```
float x;
```

```
float y;
```

```
long double nextafterl (x, y)
```

```
long double x;
```

```
long double y;
```

```
double nextafter (x, y)
```

```
double x, y;
```

```
double nexttoward (x, y)
```

```
double x;
```

```
long double y;
```

```
float nexttowardf (x, y)
```

```
float x;
```

```
long double y;
```

```
long double nexttowardl (x, y)
```

```
long double x;
```

```
long double y;
```

## Description

The **nextafterf**, **nextafterl**, and **nextafter** subroutines compute the next representable floating-point value following *x* in the direction of *y*. Thus, if *y* is less than *x*, the **nextafter** subroutine returns the largest representable floating-point number less than *x*.

The **nextafter**, **nextafterf**, and **nextafterl** subroutines return *y* if *x* equals *y*.

The **nexttoward**, **nexttowardf**, and **nexttowardl** subroutines are equivalent to the corresponding **nextafter** subroutine, except that the second parameter has type **long double** and the subroutines return *y* converted to the type of the subroutine if *x* equals *y*.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

<i>x</i>	Specifies the starting value. The next representable floating-point number is found from <i>x</i> in the direction specified by <i>y</i> .
<i>y</i>	Specifies the direction.

## Return Values

Upon successful completion, the **nextafterf**, **nextafterl**, **nextafter**, **nexttoward**, **nexttowardf**, and **nexttowardl** subroutines return the next representable floating-point value following *x* in the direction of *y*.

If *x*==*y*, *y* (of the type *x*) is returned.

If *x* is finite and the correct function value would overflow, a range error occurs and **±HUGE\_VAL**, **±HUGE\_VALF**, and **±HUGE\_VALL** (with the same sign as *x*) is returned as appropriate for the return type of the function.

If  $x$  or  $y$  is NaN, a NaN is returned.

If  $x \neq y$  and the correct subroutine value is subnormal, zero, or underflows, a range error occurs, and either the correct function value (if representable) or 0.0 is returned.

## Error Codes

For the **nextafter** subroutine, if the  $x$  parameter is finite and the correct function value would overflow, **HUGE\_VAL** is returned and **errno** is set to **ERANGE**.

## Related Information

“feclearexcept Subroutine” on page 265 and “fetestexcept Subroutine” on page 273.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## nextafterd32, nextafterd64, nextafterd128, nexttowardd32, nexttowardd64, and nexttowardd128 Subroutines

### Purpose

Compute the next representable decimal floating-point number.

### Syntax

```
#include <math.h>
_Decimal32 nextafterd32 (x, y)
_Decimal32 x;
_Decimal32 y;

_Decimal64 nextafterd64 (x, y)
_Decimal64 x;
_Decimal64 y;

_Decimal128 nextafterd128 (x, y)
_Decimal128 x;
_Decimal128 y;

_Decimal32 nexttowardd32 (x, y)
_Decimal32 x;
_Decimal128 y;

_Decimal64 nexttowardd64 (x, y)
_Decimal64 x;
_Decimal128 y;

_Decimal128 nexttowardd128 (x, y)
_Decimal128 x;
_Decimal128 y;
```

### Description

The **nextafterd32**, **nextafterd64**, and **nextafterd128** subroutines compute the next representable decimal floating-point value following the  $x$  value in the direction of the  $y$  value. Thus, if the  $y$  value is less than the  $x$  value, the **nextafter** subroutine returns the largest representable decimal floating-point number that is less than  $x$ .

If the value of  $x$  equals  $y$ , the **nextafterd32**, **nextafterd64**, and **nextafterd128** subroutines return the value of  $y$ .

The **nexttowardd32**, **nexttowardd64**, and **nexttowardd128** subroutines are equivalent to the corresponding **nextafter** subroutines, except that the second parameter has the **\_Decimal128** type, and the subroutines return the value of the *y* parameter that is converted to the type of the subroutine if the value of *x* equals that of *y*.

To check error situations, an application should set the **errno** global variable to zero and call the **feclearexcept** subroutine (**FE\_ALL\_EXCEPT**) before calling these subroutines. Upon return, if **errno** is nonzero or the **fetestexcept** subroutine (**FE\_INVALID**| **FE\_DIVBYZERO**| **FE\_OVERFLOW**| **FE\_UNDERFLOW**) is nonzero, an error has occurred.

## Parameters

**x** Specifies the starting values. The next representable decimal floating-point number is found from the **x** parameter in the direction specified by the **y** parameter.

**y** Specifies the direction.

## Return Values

Upon successful completion, the **nextafterd32**, **nextafterd64**, **nextafterd128**, **nexttowardd32**, **nexttowardd64**, and **nexttowardd128** subroutines return the next representable decimal floating-point value following the value of the **x** parameter in the direction specified by the **y** parameter.

If **x == y** (of the type **x**) is returned.

If **x** is finite and thus the correct function value overflows, a range error occurs, **±HUGE\_VAL\_D32**, **±HUGE\_VAL\_D64**, and **±HUGE\_VAL\_D128** (with the same sign as **x**) is returned respectively as appropriate for the returned type of the function.

If **x** or **y** is NaN, a NaN is returned.

If **x != y** and the correct subroutine value is subnormal, zero, or underflow, a range error occurs and either the correct function value (if representable) or a value of 0.0 is returned.

## Errors

If the value of the **x** parameter is finite and thus the correct function value overflows, a range error occurs and the **±HUGE\_VAL\_D32**, **±HUGE\_VAL\_D64**, and **±HUGE\_VAL\_D128** (with the same sign as the **x** parameter) is returned respectively according to the returned type of the function.

If the value of the **x** parameter is not equal to that of the **y** parameter, and the correct subroutine value is subnormal, zero, or underflow, a range error occurs and either the correct function value (if representable) or a value of 0.0 is returned.

## Related Information

The “**feclearexcept** Subroutine” on page 265 and the “**fetestexcept** Subroutine” on page 273.

---

## newpass Subroutine

### Purpose

Generates a new password for a user.

### Library

Security Library (**libc.a**)

## Syntax

```
#include <usersec.h>
#include <userpw.h>
```

```
char *newpass( Password)
struct userpw *Password;
```

## Description

**Note:** This subroutine has been depreciated and its use is not recommended. The “chpass Subroutine” on page 154 should be used in its place.

The **newpass** subroutine generates a new password for the user specified by the *Password* parameter. This subroutine displays a dialogue to enter and confirm the user’s new password.

Passwords can contain almost any legal value for a character but cannot contain (National Language Support (NLS) code points. Passwords cannot have more than the value specified by **MAX\_PASS**.

If a password is successfully generated, a pointer to a buffer containing the new password is returned and the last update time is reset.

**Note:** The **newpass** subroutine is not safe in a multithreaded environment. To use **newpass** in a threaded application, the application must keep the integrity of each thread.

## Parameters

*Password* Specifies a user password structure. This structure is defined in the **userpw.h** file and contains the following members:

**upw\_name**

A pointer to a character buffer containing the user name.

**upw\_passwd**

A pointer to a character buffer containing the current password.

**upw\_lastupdate**

The time the password was last changed, in seconds since the epoch.

**upw\_flags**

A bit mask containing 0 or more of the following values:

**PW\_ADMIN**

This bit indicates that password information for this user may only be changed by the root user.

**PW\_ADMCHG**

This bit indicates that the password is being changed by root and the password will have to be changed upon the next successful running of the **login** or **su** commands to this account.

## Security

**Policy: Authentication**

To change a password, the invoker must be properly authenticated.

**Note:** Programs that invoke the **newpass** subroutine should be written to conform to the authentication rules enforced by **newpass**. The **PW\_ADMCHG** flag should always be explicitly cleared unless the invoker of the command is an administrator.

## Return Values

If a new password is successfully generated, a pointer to the new encrypted password is returned. If an error occurs, a null pointer is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **newpass** subroutine fails if one or more of the following are true:

<b>EINVAL</b>	The structure passed to the <b>newpass</b> subroutine is invalid.
<b>ESAD</b>	Security authentication is denied for the invoker.
<b>EPERM</b>	The user is unable to change the password of a user with the <b>PW_ADMCHG</b> bit set, and the real user ID of the process is not the root user.
<b>ENOENT</b>	The user is not properly defined in the database.

## Implementation Specifics

This subroutine is part of Base Operating System (BOS) Runtime.

## Related Information

The “chpass Subroutine” on page 154, **getpass** (“getpass Subroutine” on page 402) subroutine, **getuserpw** (“getuserpw, putuserpw, or putuserpwhist Subroutine” on page 468) subroutine.

The **pwdadm** command.

---

## newpassx Subroutine

### Purpose

Generates a new password for a user (without a name length limit).

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
#include <userpw.h>
```

```
char *newpassx (Password)
struct userpw *Password;
```

### Description

**Note:** The **newpassx** subroutine has been obsoleted by the more current **chpassx** subroutine. Use the **chpassx** subroutine instead.

The **newpassx** subroutine generates a new password for the user specified by the *Password* parameter. The new password is then checked to ensure that it meets the password rules on the system unless the user is exempted from these restrictions. Users must have root user authority to invoke this subroutine. The password rules are defined in the **/etc/security/user** file or the administrative domain for the user and are described in both the user file and the **passwd** command.

Passwords can contain almost any legal value for a character but cannot contain National Language Support (NLS) code points. Passwords cannot have more characters than the value specified by **PASS\_MAX**.

The **newpassx** subroutine authenticates the user prior to returning the new password. If the **PW\_ADMCHG** flag is set in the **upw\_flags** member of the *Password* parameter, the supplied password is checked against the calling user's password. This is done to authenticate the user corresponding to the real user ID of the process instead of the user specified by the **upw\_name** member of the *Password* parameter structure.

If a password is successfully generated, a pointer to a buffer containing the new password is returned and the last update time is set to the current system time. The password value in the */etc/security/passwd* file or user's administrative domain is not modified.

**Note:** The **newpassx** subroutine is not safe in a multithreaded environment. To use **newpassx** in a threaded application, the application must keep the integrity of each thread.

## Parameters

*Password* Specifies a user password structure.

The fields in a **userpw** structure are defined in the **userpw.h** file, and they include the following members:

<b>upw_name</b>	Specifies the user's name.
<b>upw_passwd</b>	Specifies the user's encrypted password.
<b>upw_lastupdate</b>	Specifies the time, in seconds, since the epoch (that is, 00:00:00 GMT, 1 January 1970), when the password was last updated.
<b>upw_flags</b>	Specifies attributes of the password. This member is a bit mask of one or more of the following values, defined in the <b>userpw.h</b> file: <b>PW_NOCHECK</b> Specifies that new passwords need not meet password restrictions in effect for the system. <b>PW_ADMCHG</b> Specifies that the password was last set by an administrator and must be changed at the next successful use of the login or <b>su</b> command. <b>PW_ADMIN</b> Specifies that password information for this user can only be changed by the root user.
<b>upw_authdb</b>	Specifies the administrative domain containing the authentication data.

## Security

Policy: Authentication To change a password, the invoker must be properly authenticated.

**Note:** Programs that invoke the **newpassx** subroutine should be written to conform to the authentication rules enforced by **newpassx**. The **PW\_ADMCHG** flag should always be explicitly cleared unless the invoker of the command is an administrator.

## Return Values

If a new password is successfully generated, a pointer to the new encrypted password is returned. If an error occurs, a null pointer is returned and the **errno** global variable is set to indicate the error.



## Error Codes

The **newpassx** subroutine fails if one or more of the following is true:

<b>EINVAL</b>	The structure passed to the <b>newpassx</b> subroutine is invalid.
<b>ENOENT</b>	The user is not properly defined in the database.
<b>EPERM</b>	The user is unable to change the password of a user with the <b>PW_ADMCHG</b> bit set, and the real user ID of the process is not the root user.
<b>ESAD</b>	Security authentication is denied for the invoker.

## Related Information

The “getpass Subroutine” on page 402, “getuserpw Subroutine” on page 470.

The login Command, passwd Command, pwdadm Command.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## nftw or nftw64 Subroutine

### Purpose

Walks a file tree.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <ftw.h>
```

```
int nftw ( Path, Function, Depth, Flags)
const char *Path;
int>(*Function) ( );
int Depth;
int Flags;

int nftw64(Path,Function,Depth)
const char *Path;
int>(*Function) ( );
int Depth;
int Flags;
```

### Description

The **nftw** and **nftw64** subroutines recursively descend the directory hierarchy rooted in the *Path* parameter. The **nftw** and **nftw64** subroutines have a similar effect to **ftw** and **ftw64** except that they take an additional argument flags, which is a bitwise inclusive-OR of zero or more of the following flags:

<b>FTW_CHDIR</b>	If set, the current working directory will change to each directory as files are reported. If clear, the current working directory will not change.
<b>FTW_DEPTH</b>	If set, all files in a directory will be reported before the directory itself. If clear, the directory will be reported before any files.
<b>FTW_MOUNT</b>	If set, symbolic links will not be followed. If clear the links will be followed.
<b>FTW_PHYS</b>	If set, symbolic links will not be followed. If clear the links will be followed, and will not report the same file more than once.

For each file in the hierarchy, the **nftw** and **nftw64** subroutines call the function specified by the *Function* parameter. The **nftw** subroutine passes a pointer to a null-terminated character string containing the name of the file, a pointer to a `stat` structure containing information about the file, an integer and a pointer to an FTW structure. The **nftw64** subroutine passes a pointer to a null-terminated character string containing the name of the file, a pointer to a `stat64` structure containing information about the file, an integer and a pointer to an FTW structure.

The **nftw** subroutine uses the `stat` system call which will fail on files of size larger than 2 Gigabytes. The **nftw64** subroutine must be used if there is a possibility of files of size larger than 2 Gigabytes.

The integer passed to the *Function* parameter identifies the file type with one of the following values:

<b>FTW_F</b>	Regular file
<b>FTW_D</b>	Directory
<b>FTW_DNR</b>	Directory that cannot be read
<b>FTW_DP</b>	The <i>Object</i> is a directory and subdirectories have been visited. (This condition will only occur if <b>FTW_DEPTH</b> is included in flags).
<b>FTW_SL</b>	Symbolic Link
<b>FTW_SLN</b>	Symbolic Link that does not name an existin file. (This condition will only occur if the <b>FTW_PHYS</b> flag is not included in flags).
<b>FTW_NS</b>	File for which the <b>stat</b> structure could not be executed successfully

If the integer is **FTW\_DNR**, the files and subdirectories contained in that directory are not processed.

If the integer is **FTW\_NS**, the **stat** structure contents are meaningless. An example of a file that causes **FTW\_NS** to be passed to the *Function* parameter is a file in a directory for which you have read permission but not execute (search) permission.

The **FTW** structure pointer passed to the *Function* parameter contains `base` which is the offset of the object's filename in the pathname passed as the first argument to *Function*. The value of `level` indicates depth relative to the root of the walk.

The **nftw** and **nftw64** subroutines use one file descriptor for each level in the tree. The *Depth* parameter specifies the maximum number of file descriptors to be used. In general, the **nftw** and **nftw64** run faster if the value of the *Depth* parameter is at least as large as the number of levels in the tree. However, the value of the *Depth* parameter must not be greater than the number of file descriptors currently available for use. If the value of the *Depth* parameter is 0 or a negative number, the effect is the same as if it were 1.

Because the **nftw** and **nftw64** subroutines are recursive, it is possible for it to terminate with a memory fault due to stack overflow when applied to very deep file structures.

The **nftw** and **nftw64** subroutines use the **malloc** subroutine to allocate dynamic storage during its operation. If the **nftw** subroutine is terminated prior to its completion, such as by the **longjmp** subroutine being executed by the function specified by the *Function* parameter or by an interrupt routine, the **nftw** subroutine cannot free that storage. The storage remains allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have the function specified by the *Function* parameter return a nonzero value the next time it is called.

## Parameters

<i>Path</i>	Specifies the directory hierarchy to be searched.
<i>Function</i>	User supplied function that is called for each file encountered.
<i>Depth</i>	Specifies the maximum number of file descriptors to be used. <i>Depth</i> cannot be greater than <code>OPEN_MAX</code> which is described in the <code>sys/limits.h</code> header file.

## Return Values

If the tree is exhausted, the **nftw** and **nftw64** subroutine returns a value of 0. If the subroutine pointed to by **fn** returns a nonzero value, **nftw** and **nftw64** stops its tree traversal and returns whatever value was returned by the subroutine pointed to by **fn**. If the **nftw** and **nftw64** subroutine detects an error, it returns a -1 and sets the **errno** global variable to indicate the error.

## Error Codes

If the **nftw** or **nftw64** subroutines detect an error, a value of -1 is returned and the **errno** global variable is set to indicate the error.

The **nftw** and **nftw64** subroutine are unsuccessful if:

<b>EACCES</b>	Search permission is denied for any component of the <i>Path</i> parameter or read permission is denied for <i>Path</i> .
<b>ENAMETOOLONG</b>	The length of the path exceeds <b>PATH_MAX</b> while <b>_POSIX_NO_TRUNC</b> is in effect.
<b>ENOENT</b>	The <i>Path</i> parameter points to the name of a file that does not exist or points to an empty string.
<b>ENOTDIR</b>	A component of the <i>Path</i> parameter is not a directory.

The **nftw** subroutine is unsuccessful if:

**EOVERFLOW** A file in *Path* is of a size larger than 2 Gigabytes.

## Related Information

The **stat** or **malloc** (“**malloc**, **free**, **realloc**, **calloc**, **malloc**, **mallinfo**, **mallinfo\_heap**, **alloca**, **valloc**, or **posix\_memalign** Subroutine” on page 778) subroutine.

The **ftw** (“**ftw** or **ftw64** Subroutine” on page 324) subroutine.

---

## nl\_langinfo Subroutine

### Purpose

Returns information on the language or cultural area in a program’s locale.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <nl_types.h>
#include <langinfo.h>
```

```
char *nl_langinfo ( Item )
nl_item Item;
```

### Description

The **nl\_langinfo** subroutine returns a pointer to a string containing information relevant to the particular language or cultural area defined in the program’s locale and corresponding to the *Item* parameter. The active language or cultural area is determined by the default value of the environment variables or by the most recent call to the **setlocale** subroutine. If the **setlocale** subroutine has not been called in the program, then the default C locale values will be returned from **nl\_langinfo**.

Values for the *Item* parameter are defined in the **langinfo.h** file.

The following table summarizes the categories for which **nl\_langinfo()** returns information, the values the *Item* parameter can take, and descriptions of the returned strings. In the table, radix character refers to the character that separates whole and fractional numeric or monetary quantities. For example, a period (.) is used as the radix character in the U.S., and a comma (,) is used as the radix character in France.

Category	Value of <i>item</i>	Returned Result
LC_MONETARY	CRNCYSTR	Currency symbol and its position.
LC_NUMERIC	RADIXCHAR	Radix character.
LC_NUMERIC	THOUSEP	Separator for the thousands.
LC_MESSAGES	YESSTR	Affirmative response for yes/no queries.
LC_MESSAGES	NOSTR	Negative response for yes/no queries.
LC_TIME	D_T_FMT	String for formatting date and time.
LC_TIME	D_FMT	String for formatting date.
LC_TIME	T_FMT	String for formatting time.
LC_TIME	AM_STR	Antemeridian affix.
LC_TIME	PM_STR	Postmeridian affix.
LC_TIME	DAY_1 through DAY_7	Name of the first day of the week to the seventh day of the week.
LC_TIME	ABDAY_1 through ABDAY-7	Abbreviated name of the first day of the week to the seventh day of the week.
LC_TIME	MON_1 through MON_12	Name of the first month of the year to the twelfth month of the year.
LC_TIME	ABMON_1 through ABMON_12	Abbreviated name of the first month of the year to the twelfth month.
LC_CTYPE	CODESET	Code set currently in use in the program.

**Note:** The information returned by the **nl\_langinfo** subroutine is located in a static buffer. The contents of this buffer are overwritten in subsequent calls to the **nl\_langinfo** subroutine. Therefore, you should save the returned information.

## Parameter

*Item* Information needed from locale.

## Return Values

In a locale where language information data is not defined, the **nl\_langinfo** subroutine returns a pointer to the corresponding string in the C locale. In all locales, the **nl\_langinfo** subroutine returns a pointer to an empty string if the *Item* parameter contains an invalid setting.

The **nl\_langinfo** subroutine returns a pointer to a static area. Subsequent calls to the **nl\_langinfo** subroutine overwrite the results of a previous call.

## Related Information

The **localeconv** (“localeconv Subroutine” on page 735) subroutine, **rpmatch** subroutine, **setlocale** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Setting the Locale in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## nlist, nlist64 Subroutine

### Purpose

Gets entries from a name list.

### Library

Standard C Library (**libc.a**)

Berkeley Compatibility Library [**libbsd.a**] for the **nlist** subroutine, 32-bit programs, and POWER-based platforms

### Syntax

```
#include <nlist.h>
```

```
int nlist ( FileName, NL )
const char *FileName;
struct nlist *NL;
```

```
int nlist64 ( FileName, NL64 )
const char *FileName;
struct nlist64 *NL64;
```

### Description

The **nlist** and **nlist64** subroutines examine the name list in the object file named by the *FileName* parameter. The subroutine selectively reads a list of values and stores them into an array of **nlist** or **nlist64** structures pointed to by the *NL* or *NL64* parameter, respectively.

The name list specified by the *NL* or *NL64* parameter consists of an array of **nlist** or **nlist64** structures containing symbol names and other information. The list is terminated with an element that has a null pointer or a pointer to a null string in the **n\_name** structure member. Each symbol name is looked up in the name list of the file. If the name is found, the value of the symbol is stored in the structure, and the other fields are filled in. If the program was not compiled with the **-g** flag, the **n\_type** field may be 0.

If multiple instances of a symbol are found, the information about the last instance is stored. If a symbol is not found, all structure fields except the **n\_name** field are set to 0. Only global symbols will be found.

The **nlist** and **nlist64** subroutines run in both 32-bit and 64-bit programs that read the name list of both 32-bit and 64-bit object files, with one exception: in 32-bit programs, **nlist** will return -1 if the specified file is a 64-bit object file.

The **nlist** and **nlist64** subroutines are used to read the name list from XCOFF object files.

The **nlist64** subroutine can be used to examine the system name list kept in the kernel, by specifying **/unix** as the *FileName* parameter. The **knlist** subroutine can also be used to look up symbols in the current kernel namespace.

**Note:** The **nlist.h** header file has a *#define* field for **n\_name**. If a source file includes both **nlist.h** and **netdb.h**, there will be a conflict with the use of **n\_name**. If **netdb.h** is included after **nlist.h**, **n\_name** will be undefined. To correct this problem, **\_n\_n\_name** should be used instead. If **netdb.h** is included before **nlist.h**, and you need to refer to the **n\_name** field of *struct netent*, you should undefine **n\_name** by entering:

```
#undef n_name
```

The **nlist** subroutine in **libbsd.a** is supported only in 32-bit mode.

## Parameters

<i>FileName</i>	Specifies the name of the file containing a name list.
<i>NL</i>	Points to the array of <b>nlist</b> structures.
<i>NL64</i>	Points to the array of <b>nlist64</b> structures.

## Return Values

Upon successful completion, a 0 is returned, even if some symbols could not be found. In the **libbsd.a** version of **nlist**, the number of symbols not found in the object file's name list is returned. If the file cannot be found or if it is not a valid name list, a value of -1 is returned.

## Compatibility Interfaces

To obtain the BSD-compatible version of the subroutine 32-bit applications, compile with the **libbsd.a** library by using the **-lbsd** flag.

## Related Information

The **knlist** subroutine.

The **a.out** file in XCOFF format.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ns\_addr Subroutine

### Purpose

Address conversion routines.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
```

```
#include <netns/ns.h>
```

```
struct ns_addr(char *cp)
```

## Description

The **ns\_addr** subroutine interprets character strings representing addresses, returning binary information suitable for use in system calls.

The **ns\_addr** subroutine separates an address into one to three fields using a single delimiter and examines each field for byte separators (colon or period). The delimiters are:

- . period
- : colon
- # pound sign.

If byte separators are found, each subfield separated is taken to be a small hexadecimal number, and the entirety is taken as a network-byte-ordered quantity to be zero extended in the high-networked-order bytes. Next, the field is inspected for hyphens, which would indicate the field is a number in decimal notation with hyphens separating the millenia. The field is assumed to be a number, interpreted as hexadecimal, if a leading *0x* (as in C), a trailing *H*, (as in Mesa), or any super-octal digits are present. The field is interpreted as octal if a leading *0* is present and there are no super-octal digits. Otherwise, the field is converted as a decimal number.

## Parameter

*cp* Returns a pointer to the address of a **ns\_addr** structure.

---

## ns\_ntoa Subroutine

### Purpose

Address conversion routines.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <netns/ns.h>
```

```
char *ns_ntoa (
struct ns_addr ns)
```

### Description

The **ns\_ntoa** subroutine takes addresses and returns ASCII strings representing the address in a notation in common use in the Xerox Development Environment:

<network number> <host number> <port number>

Trailing zero fields are suppressed, and each number is printed in hexadecimal, in a format suitable for input to the **ns\_addr** subroutine. Any fields lacking super-decimal digits will have a trailing *H* appended.

**Note:** The string returned by **ns\_ntoa** resides in static memory.

### Parameter

*ns* Returns a pointer to a string.

---

## odm\_add\_obj Subroutine

### Purpose

Adds a new object into an object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_add_obj ( ClassSymbol, DataStructure)  
CLASS_SYMBOL ClassSymbol;  
struct ClassName *DataStructure;
```

### Description

The **odm\_add\_obj** subroutine takes as input the class symbol that identifies both the object class to add and a pointer to the data structure containing the object to be added.

The **odm\_add\_obj** subroutine opens and closes the object class around the subroutine if the object class was not previously opened. If the object class was previously opened, the subroutine leaves the object class open when it returns.

### Parameters

*ClassSymbol*

Specifies a class symbol identifier returned from an **odm\_open\_class** subroutine. If the **odm\_open\_class** subroutine has not been called, then this identifier is the *ClassName\_CLASS* structure that was created by the **odmcreate** command.

*DataStructure*

Specifies a pointer to an instance of the C language structure corresponding to the object class referenced by the *ClassSymbol* parameter. The structure is declared in the *.h* file created by the **odmcreate** command and has the same name as the object class.

### Return Values

Upon successful completion, an identifier for the object that was added is returned. If the **odm\_add\_obj** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

### Error Codes

Failure of the **odm\_add\_obj** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_OPEN\_ERR**
- **ODMI\_PARAMS**
- **ODMI\_READ\_ONLY**
- **ODMI\_TOOMANYCLASSES**



See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_create\_class** ("odm\_create\_class Subroutine" on page 915) subroutine, **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdnly Subroutine" on page 926) subroutine, **odm\_rm\_obj** ("odm\_rm\_obj Subroutine" on page 929) subroutine.

The **odmcreate** command.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a **.h** file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_change\_obj Subroutine

### Purpose

Changes an object in the object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_change_obj ( ClassSymbol, DataStructure)  
CLASS_SYMBOL ClassSymbol;  
struct ClassName *DataStructure;
```

### Description

The **odm\_change\_obj** subroutine takes as input the class symbol that identifies both the object class to change and a pointer to the data structure containing the object to be changed. The application program must first retrieve the object with an **odm\_get\_obj** subroutine call, change the data values in the returned structure, and then pass that structure to the **odm\_change\_obj** subroutine.

The **odm\_change\_obj** subroutine opens and closes the object class around the change if the object class was not previously opened. If the object class was previously opened, then the subroutine leaves the object class open when it returns.

### Parameters

<i>ClassSymbol</i>	Specifies a class symbol identifier returned from an <b>odm_open_class</b> subroutine. If the <b>odm_open_class</b> subroutine has not been called, then this identifier is the <i>ClassName_CLASS</i> structure that is created by the <b>odmcreate</b> command.
<i>DataStructure</i>	Specifies a pointer to an instance of the C language structure corresponding to the object class referenced by the <i>ClassSymbol</i> parameter. The structure is declared in the <b>.h</b> file created by the <b>odmcreate</b> command and has the same name as the object class.

### Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_change\_obj** subroutine fails, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_change\_obj** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_NO\_OBJECT**
- **ODMI\_OPEN\_ERR**
- **ODMI\_PARAMS**
- **ODMI\_READ\_ONLY**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_obj** ("odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine" on page 921) subroutine.

The **odmchange** command, **odmcreate** command.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a **.h** file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_close\_class Subroutine

### Purpose

Closes an ODM object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_close_class ( ClassSymbol )  
CLASS_SYMBOL ClassSymbol;
```

### Description

The **odm\_close\_class** subroutine closes the specified object class.

### Parameters

*ClassSymbol* Specifies a class symbol identifier returned from an **odm\_open\_class** subroutine. If the **odm\_open\_class** subroutine has not been called, then this identifier is the *ClassName\_CLASS* structure that was created by the **odmcreate** command.

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_close\_class** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_close\_class** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_OPEN\_ERR**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdonly Subroutine" on page 926) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_create\_class Subroutine

### Purpose

Creates an object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_create_class ( ClassSymbol )  
CLASS_SYMBOL ClassSymbol;
```

### Description

The **odm\_create\_class** subroutine creates an object class. However, the **.c** and **.h** files generated by the **odmcreate** command are required to be part of the application.

### Parameters

*ClassSymbol* Specifies a class symbol of the form *ClassName\_CLASS*, which is declared in the **.h** file created by the **odmcreate** command.

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_create\_class** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_create\_class** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_EXISTS**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_OPEN\_ERR**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_mount\_class** ("odm\_mount\_class Subroutine" on page 925) subroutine.

The **odmcreate** command.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a **.h** file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_err\_msg Subroutine

### Purpose

Returns an error message string.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_err_msg ( ODMErrno, MessageString)
long ODMErrno;
char **MessageString;
```

### Description

The **odm\_err\_msg** subroutine takes as input an *ODMErrno* parameter and an address in which to put the string pointer of the message string that corresponds to the input ODM error number. If no corresponding message is found for the input error number, a null string is returned and the subroutine is unsuccessful.

### Parameters

*ODMErrno*                      Specifies the error code for which the message string is retrieved.

*MessageString*

Specifies the address of a string pointer that will point to the returned error message string.

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_err\_msg** subroutine is unsuccessful, a value of -1 is returned, and the *MessageString* value returned is a null string.

## Examples

The following example shows the use of the **odm\_err\_msg** subroutine:

```
#include <odmi.h>
char *error_message;

...
/*-----*/
/*ODMErrno was returned from a previous ODM subroutine call.*/
/*-----*/
returnstatus = odm_err_msg ( odmerrno, &error_message );
if ( returnstatus < 0 )
    printf ( "Retrieval of error message failed\n" );
else
    printf ( error_message );
```

## Related Information

Appendix B, "ODM Error Codes," on page 1347 in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 1* describes error codes.

See Appendix B, Appendix B, "ODM Error Codes," on page 1347 for explanations of the ODM error codes.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_free\_list Subroutine

### Purpose

Frees memory previously allocated for an **odm\_get\_list** subroutine.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>

int odm_free_list ( ReturnData, DataInfo)
struct ClassName *ReturnData;
struct listinfo *DataInfo;
```

### Description

The **odm\_free\_list** subroutine recursively frees up a tree of memory object lists that were allocated for an **odm\_get\_list** subroutine.

### Parameters

*ReturnData* Points to the array of *ClassName* structures returned from the **odm\_get\_list** subroutine.

*DataInfo* Points to the **listinfo** structure that was returned from the **odm\_get\_list** subroutine. The **listinfo** structure has the following form:

```
struct listinfo {
char ClassName[16];      /* class name for query */
char criteria[256];     /* query criteria */
int num;                /* number of matches found */
int valid;              /* for ODM use */
CLASS_SYMBOL class;    /* symbol for queried class */
};
```

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_free\_list** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_free\_list** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_MAGICNO\_ERR**
- **ODMI\_PARAMS**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_list** ("odm\_get\_list Subroutine" on page 919) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_get\_by\_id Subroutine

### Purpose

Retrieves an object from an ODM object class by its ID.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
struct ClassName *odm_get_by_id( ClassSymbol, ObjectID, ReturnData)
CLASS_SYMBOL ClassSymbol;
int ObjectID;
struct ClassName *ReturnData;
```

### Description

The **odm\_get\_by\_id** subroutine retrieves an object from an object class. The object to be retrieved is specified by passing its *ObjectID* parameter from its corresponding *ClassName* structure.

### Parameters

*ClassSymbol* Specifies a class symbol identifier of the form *ClassName\_CLASS*, which is declared in the **.h** file created by the **odmcreate** command.

<i>ObjectID</i>	Specifies an identifier retrieved from the corresponding <i>ClassName</i> structure of the object class.
<i>ReturnData</i>	Specifies a pointer to an instance of the C language structure corresponding to the object class referenced by the <i>ClassSymbol</i> parameter. The structure is declared in the <i>.h</i> file created by the <b>odmcreate</b> command and has the same name as the object class.

## Return Values

Upon successful completion, a pointer to the *ClassName* structure containing the object is returned. If the **odm\_get\_by\_id** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_get\_by\_id** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_NO\_OBJECT**
- **ODMI\_OPEN\_ERR**
- **ODMI\_PARAMS**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_obj** ("odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine" on page 921), **odm\_get\_first** ("odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine" on page 921), or **odm\_get\_next** ("odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine" on page 921) subroutine.

The **odmcreate** command.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a *.h* file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_get\_list Subroutine

### Purpose

Retrieves all objects in an object class that match the specified criteria.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```

struct ClassName *odm_get_list (ClassSymbol, Criteria, ListInfo, MaxReturn, LinkDepth)
struct ClassName_CLASS ClassSymbol;
char * Criteria;
struct listinfo * ListInfo;
int MaxReturn, LinkDepth;

```

## Description

The **odm\_get\_list** subroutine takes an object class and criteria as input, and returns a list of objects that satisfy the input criteria. The subroutine opens and closes the object class around the subroutine if the object class was not previously opened. If the object class was previously opened, the subroutine leaves the object class open when it returns.

## Parameters

<i>ClassSymbol</i>	Specifies a class symbol identifier returned from an <b>odm_open_class</b> subroutine. If the <b>odm_open_class</b> subroutine has not been called, then this is the <i>ClassName_CLASS</i> structure created by the <b>odmcreate</b> command.
<i>Criteria</i>	Specifies a string that contains the qualifying criteria for selecting the objects to remove.
<i>ListInfo</i>	Specifies a structure containing information about the retrieval of the objects. The <b>listinfo</b> structure has the following form: <pre> <b>struct</b> <i>listinfo</i> { <b>char</b> <i>ClassName</i>[16];      /* class name used for query */ <b>char</b> <i>criteria</i>[256];     /* query criteria */ <b>int</b> <i>num</i>;                /* number of matches found */ <b>int</b> <i>valid</i>;              /* for ODM use */ <b>CLASS_SYMBOL</b> <i>class</i>;    /* symbol for queried class */ }; </pre>
<i>MaxReturn</i>	Specifies the expected number of objects to be returned. This is used to control the increments in which storage for structures is allocated, to reduce the <b>realloc</b> subroutine copy overhead.
<i>LinkDepth</i>	Specifies the number of levels to recurse for objects with <b>ODM_LINK</b> descriptors. A setting of 1 indicates only the top level is retrieved; 2 indicates <b>ODM_LINKs</b> will be followed from the top/first level only; 3 indicates <b>ODM_LINKs</b> will be followed at the first and second level, and so on.

## Return Values

Upon successful completion, a pointer to an array of C language structures containing the objects is returned. This structure matches that described in the **.h** file that is returned from the **odmcreate** command. If no match is found, null is returned. If the **odm\_get\_list** subroutine fails, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_get\_list** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_BAD\_CRIT**
- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INTERNAL\_ERR**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_LINK\_NOT\_FOUND**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_OPEN\_ERR**



- **ODMI\_PARAMS**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_by\_id** ("odm\_get\_by\_id Subroutine" on page 918) subroutine, **odm\_get\_obj** ("odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine") subroutine, **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdnly Subroutine" on page 926) subroutine, or **odm\_free\_list** ("odm\_free\_list Subroutine" on page 917) subroutine.

The **odmcreate** command, **odmget** command.

For information on qualifying criteria, see "Understanding ODM Object Searches" in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a **.h** file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine

### Purpose

Retrieves objects, one object at a time, from an ODM object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
struct ClassName *odm_get_obj ( ClassSymbol, Criteria, ReturnData, FIRST_NEXT)
```

```
struct ClassName *odm_get_first (ClassSymbol, Criteria, ReturnData)
```

```
struct ClassName *odm_get_next (ClassSymbol, ReturnData)
```

```
CLASS_SYMBOL ClassSymbol;
```

```
char *Criteria;
```

```
struct ClassName *ReturnData;
```

```
int FIRST_NEXT;
```

### Description

The **odm\_get\_obj**, **odm\_get\_first**, and **odm\_get\_next** subroutines retrieve objects from ODM object classes and return the objects into C language structures defined by the **.h** file produced by the **odmcreate** command.

The **odm\_get\_obj**, **odm\_get\_first**, and **odm\_get\_next** subroutines open and close the specified object class if the object class was not previously opened. If the object class was previously opened, the subroutines leave the object class open upon return.

## Parameters

<i>ClassSymbol</i>	Specifies a class symbol identifier returned from an <b>odm_open_class</b> subroutine. If the <b>odm_open_class</b> subroutine has not been called, then this identifier is the <i>ClassName</i> <b>CLASS</b> structure that was created by the <b>odmcreate</b> command.
<i>Criteria</i>	Specifies the string that contains the qualifying criteria for retrieval of the objects.
<i>ReturnData</i>	Specifies the pointer to the data structure in the <b>.h</b> file created by the <b>odmcreate</b> command. The name of the structure in the <b>.h</b> file is <i>ClassName</i> . If the <i>ReturnData</i> parameter is null ( <i>ReturnData</i> == null), space is allocated for the parameter and the calling application is responsible for freeing this space at a later time.  If variable length character strings (vchar) are returned, they are referenced by pointers in the <i>ReturnData</i> structure. Calling applications must free each vchar between each call to the <b>odm_get</b> subroutines; otherwise storage will be lost.
<i>FIRST_NEXT</i>	Specifies whether to get the first object that matches the criteria or the next object. Valid values are:  <b>ODM_FIRST</b> Retrieve the first object that matches the search criteria.  <b>ODM_NEXT</b> Retrieve the next object that matches the search criteria. The <i>Criteria</i> parameter is ignored if the <i>FIRST_NEXT</i> parameter is set to <b>ODM_NEXT</b> .

## Return Values

Upon successful completion, a pointer to the retrieved object is returned. If no match is found, null is returned. If an **odm\_get\_obj**, **odm\_get\_first**, or **odm\_get\_next** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_get\_obj**, **odm\_get\_first** or **odm\_get\_next** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_BAD\_CRIT**
- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INTERNAL\_ERR**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_OPEN\_ERR**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_list** ("odm\_get\_list Subroutine" on page 919) subroutine, **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdonly Subroutine" on page 926) subroutine, **odm\_rm\_by\_id** ("odm\_rm\_by\_id Subroutine" on page 927) subroutine, **odm\_rm\_obj** ("odm\_rm\_obj Subroutine" on page 929) subroutine.

The **odmcreate** command, **odmget** command.

For more information about qualifying criteria, see "Understanding ODM Object Searches" in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a `.h` file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_initialize Subroutine

### Purpose

Prepares ODM for use by an application.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
int odm_initialize( )
```

### Description

The **odm\_initialize** subroutine starts ODM for use with an application program.

### Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_initialize** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

### Error Codes

Failure of the **odm\_initialize** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_INVALID\_PATH**
- **ODMI\_MALLOC\_ERR**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

### Related Information

The **odm\_terminate** ("odm\_terminate Subroutine" on page 933) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_lock Subroutine

### Purpose

Puts an exclusive lock on the requested path name.

### Library

Object Data Manager Library (**libodm.a**)

## Syntax

```
#include <odmi.h>
```

```
int odm_lock ( LockPath, Timeout)
char *LockPath;
int Timeout;
```

## Description

The **odm\_lock** subroutine is used by an application to prevent other applications or methods from accessing an object class or group of object classes. A lock on a directory path name does not prevent another application from acquiring a lock on a subdirectory or object class within that directory.

**Note:** Coordination of locking is the responsibility of the application accessing the object classes.

The **odm\_lock** subroutine returns a lock identifier that is used to call the **odm\_unlock** subroutine.

## Parameters

<i>LockPath</i>	Specifies a string containing the path name in the file system in which to locate object classes or the path name to an object class to lock.
<i>Timeout</i>	Specifies the amount of time, in seconds, to wait if another application or method holds a lock on the requested object class or classes. The possible values for the <i>Timeout</i> parameter are:  <i>Timeout</i> = <b>ODM_NOWAIT</b> The <b>odm_lock</b> subroutine is unsuccessful if the lock cannot be granted immediately.  <i>Timeout</i> = <i>Integer</i> The <b>odm_lock</b> subroutine waits the specified amount of seconds to retry an unsuccessful lock request.  <i>Timeout</i> = <b>ODM_WAIT</b> The <b>odm_lock</b> subroutine waits until the locked path name is freed from its current lock and then locks it.

## Return Values

Upon successful completion, a lock identifier is returned. If the **odm\_lock** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_lock** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_BAD\_LOCK**
- **ODMI\_BAD\_TIMEOUT**
- **ODMI\_BAD\_TOKEN**
- **ODMI\_LOCK\_BLOCKED**
- **ODMI\_LOCK\_ENV**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_UNLOCK**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_unlock** ("odm\_unlock Subroutine" on page 934) subroutine.

## odm\_mount\_class Subroutine

### Purpose

Retrieves the class symbol structure for the specified object class name.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
CLASS_SYMBOL odm_mount_class ( ClassName)  
char *ClassName;
```

### Description

The **odm\_mount\_class** subroutine retrieves the class symbol structure for a specified object class. The subroutine can be called by applications (for example, the ODM commands) that have no previous knowledge of the structure of an object class before trying to access that class. The **odm\_mount\_class** subroutine determines the class description from the object class header information and creates a **CLASS\_SYMBOL** object class that is returned to the caller.

The object class is not opened by the **odm\_mount\_class** subroutine. Calling the subroutine subsequent times for an object class that is already open or mounted returns the same **CLASS\_SYMBOL** object class.

Mounting a class that links to another object class recursively mounts to the linked class. However, if the recursive mount is unsuccessful, the original **odm\_mount\_class** subroutine does not fail; the **CLASS\_SYMBOL** object class is set up with a null link.

### Parameters

*ClassName*        Specifies the name of an object class from which to retrieve the class description.

### Return Values

Upon successful completion, a **CLASS\_SYMBOL** is returned. If the **odm\_mount\_class** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

### Error Codes

Failure of the **odm\_mount\_class** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_BAD\_CLASSNAME**
- **ODMI\_BAD\_CLXNNAME**
- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_CLXNMAGICNO\_ERR**
- **ODMI\_INVALID\_CLASS**
- **ODMI\_INVALID\_CLXN**

- ODMI\_MAGICNO\_ERR
- ODMI\_MALLOC\_ERR
- ODMI\_OPEN\_ERR
- ODMI\_PARAMS
- ODMI\_TOOMANYCLASSES
- ODMI\_TOOMANYCLASSES

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_create\_class** ("odm\_create\_class Subroutine" on page 915) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_open\_class or odm\_open\_class\_rdonly Subroutine

### Purpose

Opens an ODM object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
CLASS_SYMBOL odm_open_class ( ClassSymbol )
CLASS_SYMBOL ClassSymbol ;
```

```
CLASS_SYMBOL odm_open_class_rdonly ( ClassSymbol )
CLASS_SYMBOL ClassSymbol ;
```

### Description

The **odm\_open\_class** subroutine can be called to open an object class. Most subroutines implicitly open a class if the class is not already open. However, an application may find it useful to perform an explicit open if, for example, several operations must be done on one object class before closing the class. The **odm\_open\_class\_rdonly** subroutine opens an **odm** database in read-only mode.

### Parameter

<i>ClassSymbol</i>	Specifies a class symbol of the form <i>ClassName_CLASS</i> that is declared in the <b>.h</b> file created by the <b>odmcreate</b> command.
--------------------	---

### Return Values

Upon successful completion, a *ClassSymbol* parameter for the object class is returned. If the **odm\_open\_class** or **odm\_open\_class\_rdonly** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the `odm_open_class` or `odm_open_class_rdonly` subroutine sets the `odmerrno` variable to one of the following error codes:

- `ODMI_CLASS_DNE`
- `ODMI_CLASS_PERMS`
- `ODMI_INVALID_PATH`
- `ODMI_MAGICNO_ERR`
- `ODMI_OPEN_ERR`
- `ODMI_TOOMANYCLASSES`

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The `odm_close_class` ("odm\_close\_class Subroutine" on page 914) subroutine.

The `odmcreate` command.

See ODM Example Code and Output in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs* for an example of a `.h` file.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_rm\_by\_id Subroutine

### Purpose

Removes objects specified by their IDs from an ODM object class.

### Library

Object Data Manager Library (`libodm.a`)

### Syntax

```
#include <odmi.h>
```

```
int odm_rm_by_id( ClassSymbol, ObjectID)
CLASS_SYMBOL ClassSymbol;
int ObjectID;
```

### Description

The `odm_rm_by_id` subroutine is called to delete an object from an object class. The object to be deleted is specified by passing its object ID from its corresponding `ClassName` structure.

### Parameters

<i>ClassSymbol</i>	Identifies a class symbol returned from an <code>odm_open_class</code> subroutine. If the <code>odm_open_class</code> subroutine has not been called, this is the <code>ClassName_CLASS</code> structure that was created by the <code>odmcreate</code> command.
<i>ObjectID</i>	Identifies the object. This information is retrieved from the corresponding <code>ClassName</code> structure of the object class.

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_rm\_by\_id** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_rm\_by\_id** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_FORK**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_NO\_OBJECT**
- **ODMI\_OPEN\_ERR**
- **ODMI\_OPEN\_PIPE**
- **ODMI\_PARAMS**
- **ODMI\_READ\_ONLY**
- **ODMI\_READ\_PIPE**
- **ODMI\_TOOMANYCLASSES**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_obj** ("odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine" on page 921) subroutine, **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdonly Subroutine" on page 926) subroutine.

The **odmdelete** command.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_rm\_class Subroutine

### Purpose

Removes an object class from the file system.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_rm_class ( ClassSymbol)  
CLASS_SYMBOL ClassSymbol;
```



## Description

The **odm\_rm\_class** subroutine removes an object class from the file system. All objects in the specified class are deleted.

## Parameter

*ClassSymbol* Identifies a class symbol returned from the **odm\_open\_class** subroutine. If the **odm\_open\_class** subroutine has not been called, this is the *ClassName\_CLASS* structure created by the **odmcreate** command.

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_rm\_class** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_rm\_class** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_OPEN\_ERR**
- **ODMI\_TOOMANYCLASSES**
- **ODMI\_UNLINKCLASS\_ERR**
- **ODMI\_UNLINKCLXN\_ERR**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdonly Subroutine" on page 926) subroutine.

The **odmcreate** command, **odmdrop** command.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_rm\_obj Subroutine

### Purpose

Removes objects from an ODM object class.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_rm_obj ( ClassSymbol, Criteria)
CLASS_SYMBOL ClassSymbol;
char *Criteria;
```

## Description

The **odm\_rm\_obj** subroutine deletes objects from an object class.

## Parameters

*ClassSymbol* Identifies a class symbol returned from an **odm\_open\_class** subroutine. If the **odm\_open\_class** subroutine has not been called, this is the *ClassName\_CLASS* structure that was created by the **odmcreate** command.

*Criteria* Contains as a string the qualifying criteria for selecting the objects to remove.

## Return Values

Upon successful completion, the number of objects deleted is returned. If the **odm\_rm\_obj** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_rm\_obj** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_BAD\_CRIT**
- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_FORK**
- **ODMI\_INTERNAL\_ERR**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_OPEN\_ERR**
- **ODMI\_OPEN\_PIPE**
- **ODMI\_PARAMS**
- **ODMI\_READ\_ONLY**
- **ODMI\_READ\_PIPE**
- **ODMI\_TOOMANYCLASSES**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_add\_obj** ("odm\_add\_obj Subroutine" on page 912) subroutine, **odm\_open\_class** ("odm\_open\_class or odm\_open\_class\_rdonly Subroutine" on page 926) subroutine.

The **odmcreate** command, **odmdelete** command.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

For information on qualifying criteria, see "Understanding ODM Object Searches" in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_run\_method Subroutine

### Purpose

Runs a specified method.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_run_method(MethodName, MethodParameters, NewStdOut, NewStdError)
char * MethodName, * MethodParameters;
char ** NewStdOut, ** NewStdError;
```

### Description

The **odm\_run\_method** subroutine takes as input the name of the method to run, any parameters for the method, and addresses of locations for the **odm\_run\_method** subroutine to store pointers to the stdout (standard output) and stderr (standard error output) buffers. The application uses the pointers to access the stdout and stderr information generated by the method.

### Parameters

<i>MethodName</i>	Indicates the method to execute. The method can already be known by the applications, or can be retrieved as part of an <b>odm_get_obj</b> subroutine call.
<i>MethodParameters</i>	Specifies a list of parameters for the specified method.
<i>NewStdOut</i>	Specifies the address of a pointer to the memory where the standard output of the method is stored. If the <i>NewStdOut</i> parameter points to a null value ( <i>*NewStdOut</i> == NULL), standard output is not captured.
<i>NewStdError</i>	Specifies the address of a pointer to the memory where the standard error output of the method will be stored. If the <i>NewStdError</i> parameter points to a null value ( <i>*NewStdError</i> == NULL), standard error output is not captured.

### Return Values

If successful, the **odm\_run\_method** subroutine returns the exit status and *out\_ptr* and *err\_ptr* should contain the relevant information. If unsuccessful, the **odm\_run\_method** subroutine will return -1 and set the **odmerrno** variable to an error code.

**Note:** AIX methods usually return the exit code defined in the **cf.h** file if the methods exit on error.

### Error Codes

Failure of the **odm\_run\_method** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_FORK**
- **ODMI\_MALLOC\_ERR**
- **ODMI\_OPEN\_PIPE**
- **ODMI\_PARAMS**
- **ODMI\_READ\_PIPE**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_get\_obj** (“odm\_get\_obj, odm\_get\_first, or odm\_get\_next Subroutine” on page 921) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_set\_path Subroutine

### Purpose

Sets the default path for locating object classes.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
char *odm_set_path ( NewPath)  
char *NewPath;
```

### Description

The **odm\_set\_path** subroutine is used to set the default path for locating object classes. The subroutine allocates memory, sets the default path, and returns the pointer to memory. Once the operation is complete, the calling application should free the pointer using the **free** (“malloc, free, realloc, calloc, malloc, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778) subroutine.

### Parameters

*NewPath*      Contains, as a string, the path name in the file system in which to locate object classes.

### Return Values

Upon successful completion, a string pointing to the previous default path is returned. If the **odm\_set\_path** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

### Error Codes

Failure of the **odm\_set\_path** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_INVALID\_PATH**
- **ODMI\_MALLOC\_ERR**

See Appendix B, “ODM Error Codes” for explanations of the ODM error codes.

## Related Information

The **free** (“malloc, free, realloc, calloc, malloc, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_set\_perms Subroutine

### Purpose

Sets the default permissions for an ODM object class at creation time.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>
```

```
int odm_set_perms ( NewPermissions )  
int NewPermissions;
```

### Description

The **odm\_set\_perms** subroutine defines the default permissions to assign to object classes at creation.

### Parameters

*NewPermissions*                      Specifies the new default permissions parameter as an integer.

### Return Values

Upon successful completion, the current default permissions are returned. If the **odm\_set\_perms** subroutine is unsuccessful, a value of -1 is returned.

### Related Information

See Appendix B, "ODM Error Codes," on page 1347 for explanations of the ODM error codes.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_terminate Subroutine

### Purpose

Terminates an ODM session.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>  
int odm_terminate ( )
```

### Description

The **odm\_terminate** subroutine performs the cleanup necessary to terminate an ODM session. After running an **odm\_terminate** subroutine, an application must issue an **odm\_initialize** subroutine to resume ODM operations.

## Return Values

Upon successful completion, a value of 0 is returned. If the **odm\_terminate** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_terminate** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_CLASS\_DNE**
- **ODMI\_CLASS\_PERMS**
- **ODMI\_INVALID\_CLXN**
- **ODMI\_INVALID\_PATH**
- **ODMI\_LOCK\_ID**
- **ODMI\_MAGICNO\_ERR**
- **ODMI\_OPEN\_ERR**
- **ODMI\_TOOMANYCLASSES**
- **ODMI\_UNLOCK**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_initialize** ("odm\_initialize Subroutine" on page 923) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## odm\_unlock Subroutine

### Purpose

Releases a lock put on a path name.

### Library

Object Data Manager Library (**libodm.a**)

### Syntax

```
#include <odmi.h>

int odm_unlock ( LockID)
int LockID;
```

### Description

The **odm\_unlock** subroutine releases a previously granted lock on a path name. This path name can be a directory containing subdirectories and object classes.

### Parameters

*LockID* Identifies the lock returned from the **odm\_lock** subroutine.

## Return Values

Upon successful completion a value of 0 is returned. If the **odm\_unlock** subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to an error code.

## Error Codes

Failure of the **odm\_unlock** subroutine sets the **odmerrno** variable to one of the following error codes:

- **ODMI\_LOCK\_ID**
- **ODMI\_UNLOCK**

See Appendix B, "ODM Error Codes" for explanations of the ODM error codes.

## Related Information

The **odm\_lock** ("odm\_lock Subroutine" on page 923) subroutine.

Object Data Manager (ODM) Overview for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## open, openx, open64, creat, or creat64 Subroutine

### Purpose

Opens a file for reading or writing.

### Syntax

```
#include <fcntl.h>
```

```
int open ( Path, OFlag [, Mode])
const char *Path;
int OFlag;
mode_t Mode;
```

```
int openx (Path, OFlag, Mode, Extension)
const char *Path;
int OFlag;
mode_t Mode;
int Extension;
```

```
int creat (Path, Mode)
const char *Path;
mode_t Mode;
```

**Note:** The **open64** and **creat64** subroutines apply to AIX 4.2 and later releases.

```
int open64 (Path, OFlag [, Mode])
const char *Path;
int OFlag;
mode_t Mode;
```

```
int creat64 (Path, Mode)
const char *Path;
mode_t Mode;
```

### Description

**Note:** The **open64** and **creat64** subroutines apply to AIX 4.2 and later releases.

The **open**, **openx**, and **creat** subroutines establish a connection between the file named by the *Path* parameter and a file descriptor. The opened file descriptor is used by subsequent I/O subroutines, such as **read** and **write**, to access that file.

The **openx** subroutine is the same as the **open** subroutine, with the addition of an *Extension* parameter, which is provided for device driver use. The **creat** subroutine is equivalent to the **open** subroutine with the **O\_WRONLY**, **O\_CREAT**, and **O\_TRUNC** flags set.

The returned file descriptor is the lowest file descriptor not previously open for that process. No process can have more than **OPEN\_MAX** file descriptors open simultaneously.

The file offset, marking the current position within the file, is set to the beginning of the file. The new file descriptor is set to remain open across exec (“exec: **execl**, **execle**, **execlp**, **execv**, **execve**, **execvp**, or **exec** Subroutine” on page 236) subroutines.

The **open64** and **creat64** subroutines are equivalent to the **open** and **creat** subroutines except that the **O\_LARGEFILE** flag is set in the open file description associated with the returned file descriptor. This flag allows files larger than **OFF\_MAX** to be accessed. If the caller attempts to open a file larger than **OFF\_MAX** and **O\_LARGEFILE** is not set, the open will fail and **errno** will be set to **E\_OVERFLOW**.

In the large file enabled programming environment, **open** is redefined to be **open64** and **creat** is redefined to be **creat64**.

## Parameters

*Path*                      Specifies the file to be opened.



*Mode*

Specifies the read, write, and execute permissions of the file to be created (requested by the **O\_CREAT** flag). If the file already exists, this parameter is ignored. The *Mode* parameter is constructed by logically ORing one or more of the following values, which are defined in the **sys/mode.h** file:

**S\_ISUID**

Enables the **setuid** attribute for an executable file. A process executing this program acquires the access rights of the owner of the file.

**S\_ISGID**

Enables the **setgid** attribute for an executable file. A process executing this program acquires the access rights of the group of the file. Also, enables the group-inheritance attribute for a directory. Files created in this directory have a group equal to the group of the directory.

The following attributes apply only to files that are directly executable. They have no meaning when applied to executable text files such as shell scripts and **awk** scripts.

**S\_ISVTX**

Enables the **link/unlink** attribute for a directory. Files cannot be linked to in this directory. Files can only be unlinked if the requesting process has write permission for the directory and is either the owner of the file or the directory.

**S\_ISVTX**

Enables the **save text** attribute for an executable file. The program is not unmapped after usage.

**S\_ENFMT**

Enables enforcement-mode record locking for a regular file. File locks requested with the **lockf** subroutine are enforced.

**S\_IRUSR**

Permits the file's owner to read it.

**S\_IWUSR**

Permits the file's owner to write to it.

**S\_IXUSR**

Permits the file's owner to execute it (or to search the directory).

**S\_IRGRP**

Permits the file's group to read it.

**S\_IWGRP**

Permits the file's group to write to it.

**S\_IXGRP**

Permits the file's group to execute it (or to search the directory).

**S\_IROTH**

Permits others to read the file.

**S\_IWOTH**

Permits others to write to the file.

**S\_IXOTH**

Permits others to execute the file (or to search the directory).

Other mode values exist that can be set with the **mknod** subroutine but not with the **chmod** subroutine.

*Extension*

Provides communication with character device drivers that require additional information or return additional status. Each driver interprets the *Extension* parameter in a device-dependent way, either as a value or as a pointer to a communication area. Drivers must apply reasonable defaults when the *Extension* parameter value is 0.

*OFlag*

Specifies the type of access, special open processing, the type of update, and the initial state of the open file. The parameter value is constructed by logically ORing special open processing flags. These flags are defined in the **fcntl.h** file and are described in the following flags.

## Flags That Specify Access Type

The following *OFlag* parameter flag values specify type of access:

<b>O_RDONLY</b>	The file is opened for reading only.
<b>O_WRONLY</b>	The file is opened for writing only.
<b>O_RDWR</b>	The file is opened for both reading and writing.

**Note:** One of the file access values must be specified. Do not use **O\_RDONLY**, **O\_WRONLY**, or **O\_RDWR** together. If none is set, none is used, and the results are unpredictable.

## Flags That Specify Special Open Processing

The following *OFlag* parameter flag values specify special open processing:

<b>O_CREAT</b>	<p>If the file exists, this flag has no effect, except as noted under the <b>O_EXCL</b> flag. If the file does not exist, a regular file is created with the following characteristics:</p> <ul style="list-style-type: none"><li>• The owner ID of the file is set to the effective user ID of the process.</li><li>• The group ID of the file is set to the group ID of the parent directory if the parent directory has the <b>SetGroupID</b> attribute (<b>S_ISGID</b> bit) set. Otherwise, the group ID of the file is set to the effective group ID of the calling process.</li><li>• The file permission and attribute bits are set to the value of the <i>Mode</i> parameter, modified as follows:<ul style="list-style-type: none"><li>– All bits set in the process file mode creation mask are cleared. (The file creation mask is described in the <b>umask</b> subroutine.)</li><li>– The <b>S_ISVTX</b> attribute bit is cleared.</li></ul></li></ul>
<b>O_EXCL</b>	<p>If the <b>O_EXCL</b> and <b>O_CREAT</b> flags are set, the open is unsuccessful if the file exists.</p> <p><b>Note:</b> The <b>O_EXCL</b> flag is not fully supported for Network File Systems (NFS). The NFS protocol does not guarantee the designed function of the <b>O_EXCL</b> flag.</p>
<b>O_NSHARE</b>	<p>Assures that no process has this file open and precludes subsequent opens. If the file is on a physical file system and is already open, this open is unsuccessful and returns immediately unless the <i>OFlag</i> parameter also specifies the <b>O_DELAY</b> flag. This flag is effective only with physical file systems.</p> <p><b>Note:</b> This flag is not supported by NFS.</p>
<b>O_RSHARE</b>	<p>Assures that no process has this file open for writing and precludes subsequent opens for writing. The calling process can request write access. If the file is on a physical file system and is open for writing or open with the <b>O_NSHARE</b> flag, this open fails and returns immediately unless the <i>OFlag</i> parameter also specifies the <b>O_DELAY</b> flag.</p> <p><b>Note:</b> This flag is not supported by NFS.</p>
<b>O_DEFER</b>	<p>The file is opened for deferred update. Changes to the file are not reflected on permanent storage until an <b>fsync</b> (“<b>fsync</b> or <b>fsync_range</b> Subroutine” on page 321) subroutine operation is performed. If no <b>fsync</b> subroutine operation is performed, the changes are discarded when the file is closed.</p> <p><b>Note:</b> This flag is not supported by NFS or JFS2, and the flag will be quietly ignored.</p> <p><b>Note:</b> This flag causes modified pages to be backed by paging space. Before using this flag make sure there is sufficient paging space.</p>
<b>O_NOCTTY</b>	<p>This flag specifies that the controlling terminal should not be assigned during this open.</p>
<b>O_TRUNC</b>	<p>If the file does not exist, this flag has no effect. If the file exists, is a regular file, and is successfully opened with the <b>O_RDWR</b> flag or the <b>O_WRONLY</b> flag, all of the following apply:</p> <ul style="list-style-type: none"><li>• The length of the file is truncated to 0.</li><li>• The owner and group of the file are unchanged.</li><li>• The <b>SetUserID</b> attribute of the file mode is cleared.</li><li>• The <b>SetUserID</b> attribute of the file is cleared.</li></ul>
<b>O_DIRECT</b>	<p>This flag specifies that direct i/o will be used for this file while it is opened.</p>

**O\_CIO** This flag specifies that concurrent i/o (CIO) will be used for the file while it is opened. Because implementing concurrent readers and writers utilizes the direct I/O path (with more specific requirements to improve performance for running database on file system), this flag will override the **O\_DIRECT** flag if the two options are specified at the same time. Currently, only JFS2 and namefs, which includes a selected subset of JFS2 files/directories, support CIO. The length of data to be read/written and the file offset must be page-aligned to be transferred as direct i/o with concurrent readers and writers.

The **O\_CIO** flag is exclusive. If the file is opened in any other way (for example, using the **O\_DIO** flag or opening the file normally), the open will fail. If the file is opened using the **O\_CIO** and another process attempts to open the file another way, the open will fail. The **O\_CIO** flag also prevents the **mmap** subroutine and the **shmat** subroutine access to the file. The **mmap** subroutine and the **shmat** subroutine return **EINVAL** if they are used on a file that was opened using the **O\_CIO** flag.

**O\_SNAPSHOT** The file being opened contains a JFS2 snapshot. Subsequent **read** calls using this file descriptor will read the cooked snapshot rather than the raw snapshot blocks. A snapshot can only have one active open file descriptor for it.

The **open** subroutine is unsuccessful if any of the following conditions are true:

- The file supports enforced record locks and another process has locked a portion of the file.
- The file is on a physical file system and is already open with the **O\_RSHARE** flag or the **O\_NSHARE** flag.
- The file does not allow write access.
- The file is already opened for deferred update.

### Flag That Specifies Type of Update

A program can request some control on when updates should be made permanent for a regular file opened for write access. The following *OFlag* parameter values specify the type of update performed:

**O\_SYNC:** If set, updates to regular files and writes to block devices are synchronous updates. File update is performed by the following subroutines:

- **fclear**
- **ftruncate**
- **open** with **O\_TRUNC**
- **write**

On return from a subroutine that performs a synchronous update (any of the preceding subroutines, when the **O\_SYNC** flag is set), the program is assured that all data for the file has been written to permanent storage, even if the file is also open for deferred update.

**Note:** The **O\_DSYNC** flag applies to AIX 4.2.1 and later releases.

**O\_DSYNC:** If set, the file data as well as all file system meta-data required to retrieve the file data are written to their permanent storage locations. File attributes such as access or modification times are not required to retrieve file data, and as such, they are not guaranteed to be written to their permanent storage locations before the preceding subroutines return. (Subroutines listed in the **O\_SYNC** description.)

**O\_SYNC | O\_DSYNC:** If both flags are set, the file's data and all of the file's meta-data (including access time) are written to their permanent storage locations.

**Note:** The **O\_RSYNC** flag applies to AIX 4.3 and later releases.

**O\_RSYNC:** This flag is used in combination with **O\_SYNC** or **D\_SYNC**, and it extends their write operation behaviors to read operations. For example, when **O\_SYNC** and **R\_SYNC** are both set, a read operation will not return until the file's data and all of the file's meta-data (including access time) are written to their permanent storage locations.

## Flags That Define the Open File Initial State

The following *OFlag* parameter flag values define the initial state of the open file:

<b>O_APPEND</b>	The file pointer is set to the end of the file prior to each write operation.
<b>O_DELAY</b>	Specifies that if the <b>open</b> subroutine could not succeed due to an inability to grant the access on a physical file system required by the <b>O_RSHARE</b> flag or the <b>O_NSHARE</b> flag, the process blocks instead of returning the <b>ETXTBSY</b> error code.
<b>O_NDELAY</b>	Opens with no delay.
<b>O_NONBLOCK</b>	Specifies that the <b>open</b> subroutine should not block.

The **O\_NDELAY** flag and the **O\_NONBLOCK** flag are identical except for the value returned by the **read** and **write** subroutines. These flags mean the process does not block on the state of an object, but does block on input or output to a regular file or block device.

The **O\_DELAY** flag is relevant only when used with the **O\_NSHARE** or **O\_RSHARE** flags. It is unrelated to the **O\_NDELAY** and **O\_NONBLOCK** flags.

## General Notes on OFlag Parameter Flags

The effect of the **O\_CREAT** flag is immediate, even if the file is opened with the **O\_DEFER** flag.

When opening a file on a physical file system with the **O\_NSHARE** flag or the **O\_RSHARE** flag, if the file is already open with conflicting access the following can occur:

- If the **O\_DELAY** flag is clear (the default), the **open** subroutine is unsuccessful.
- If the **O\_DELAY** flag is set, the **open** subroutine blocks until there is no conflicting open. There is no deadlock detection for processes using the **O\_DELAY** flag.

When opening a file on a physical file system that has already been opened with the **O\_NSHARE** flag, the following can occur:

- If the **O\_DELAY** flag is clear (the default), the open is unsuccessful immediately.
- If the **O\_DELAY** flag is set, the open blocks until there is no conflicting open.

When opening a file with the **O\_RDWR**, **O\_WRONLY**, or **O\_TRUNC** flag, and the file is already open with the **O\_RSHARE** flag:

- If the **O\_DELAY** flag is clear (the default), the open is unsuccessful immediately.
- If the **O\_DELAY** flag is set, the open blocks until there is no conflicting open.

When opening a first-in-first-out (FIFO) with the **O\_RDONLY** flag, the following can occur:

- If the **O\_NDELAY** and **O\_NONBLOCK** flags are clear, the open blocks until a process opens the file for writing. If the file is already open for writing (even by the calling process), the **open** subroutine returns without delay.
- If the **O\_NDELAY** flag or the **O\_NONBLOCK** flag is set, the open succeeds immediately even if no process has the FIFO open for writing.

When opening a FIFO with the **O\_WRONLY** flag, the following can occur:

- If the **O\_NDELAY** and **O\_NONBLOCK** flags are clear (the default), the open blocks until a process opens the file for reading. If the file is already open for writing (even by the calling process), the **open** subroutine returns without delay.
- If the **O\_NDELAY** flag or the **O\_NONBLOCK** flag is set, the **open** subroutine returns an error if no process currently has the file open for reading.

When opening a block special or character special file that supports nonblocking opens, such as a terminal device, the following can occur:

- If the **O\_NDELAY** and **O\_NONBLOCK** flags are clear (the default), the open blocks until the device is ready or available.
- If the **O\_NDELAY** flag or the **O\_NONBLOCK** flag is set, the **open** subroutine returns without waiting for the device to be ready or available. Subsequent behavior of the device is device-specific.

Any additional information on the effect, if any, of the **O\_NDELAY**, **O\_RSHARE**, **O\_NSHARE**, and **O\_DELAY** flags on a specific device is documented in the description of the special file related to the device type.

If path refers to a STREAMS file, *oflag* may be constructed from **O\_NONBLOCK** OR-ed with either **O\_RDONLY**, **O\_WRONLY** or **O\_RDWR**. Other flag values are not applicable to STREAMS devices and have no effect on them. The value **O\_NONBLOCK** affects the operation of STREAMS drivers and certain functions applied to file descriptors associated with STREAMS files. For STREAMS drivers, the implementation of **O\_NONBLOCK** is device-specific.

If path names the master side of a pseudo-terminal device, then it is unspecified whether **open** locks the slave side so that it cannot be opened. Portable applications must call **unlockpt** before opening the slave side.

The largest value that can be represented correctly in an object of type **off\_t** will be established as the offset maximum in the open file description.

## Return Values

Upon successful completion, the file descriptor, a nonnegative integer, is returned. Otherwise, a value of -1 is returned, no files are created or modified, and the **errno** global variable is set to indicate the error.

## Error Codes

The **open**, **openx**, and **creat** subroutines are unsuccessful and the named file is not opened if one or more of the following are true:

<b>EACCES</b>	One of the following is true: <ul style="list-style-type: none"> <li>• The file exists and the type of access specified by the <i>OFlag</i> parameter is denied.</li> <li>• Search permission is denied on a component of the path prefix specified by the <i>Path</i> parameter. Access could be denied due to a secure mount.</li> <li>• The file does not exist and write permission is denied for the parent directory of the file to be created.</li> <li>• The <b>O_TRUNC</b> flag is specified and write permission is denied.</li> </ul>
<b>EAGAIN</b>	The <b>O_TRUNC</b> flag is set and the named file contains a record lock owned by another process.
<b>EDQUOT</b>	The directory in which the entry for the new link is being placed cannot be extended, or an i-node could not be allocated for the file, because the user or group quota of disk blocks or i-nodes in the file system containing the directory has been exhausted.
<b>EEXIST</b>	The <b>O_CREAT</b> and <b>O_EXCL</b> flags are set and the named file exists.

<b>EFBIG</b>	An attempt was made to write a file that exceeds the process' file limit or the maximum file size. If the user has set the environment variable <b>XPG_SUS_ENV=ON</b> prior to execution of the process, then the <b>SIGXFSZ</b> signal is posted to the process when exceeding the process' file size limit.
<b>EINTR</b>	A signal was caught during the <b>open</b> subroutine.
<b>EIO</b>	The <i>path</i> parameter names a STREAMS file and a hangup or error occurred.
<b>EISDIR</b>	Named file is a directory and write access is required (the <b>O_WRONLY</b> or <b>O_RDWR</b> flag is set in the <i>OFlag</i> parameter).
<b>EMFILE</b>	The system limit for open file descriptors per process has already been reached ( <b>OPEN_MAX</b> ).
<b>ENAMETOOLONG</b>	The length of the <i>Path</i> parameter exceeds the system limit ( <b>PATH_MAX</b> ); or a path-name component is longer than <b>NAME_MAX</b> and <b>_POSIX_NO_TRUNC</b> is in effect.
<b>ENFILE</b>	The system file table is full.
<b>ENOENT</b>	The <b>O_CREAT</b> flag is not set and the named file does not exist; or the <b>O_CREAT</b> flag is not set and either the path prefix does not exist or the <i>Path</i> parameter points to an empty string.
<b>ENOMEM</b>	The <i>Path</i> parameter names a STREAMS file and the system is unable to allocate resources.
<b>ENOSPC</b>	The directory or file system that would contain the new file cannot be extended.
<b>ENOSR</b>	The <i>Path</i> argument names a STREAMS-based file and the system is unable to allocate a STREAM.
<b>ENOTDIR</b>	A component of the path prefix specified by the <i>Path</i> component is not a directory.
<b>ENXIO</b>	One of the following is true: <ul style="list-style-type: none"> <li>• Named file is a character special or block special file, and the device associated with this special file does not exist.</li> <li>• Named file is a multiplexed special file and either the channel number is outside of the valid range or no more channels are available.</li> <li>• The <b>O_DELAY</b> flag or the <b>O_NONBLOCK</b> flag is set, the named file is a FIFO, the <b>O_WRONLY</b> flag is set, and no process has the file open for reading.</li> </ul>
<b>E_OVERFLOW</b>	A file greater than one terabyte was opened on the 32-bit kernel in JFS2. The exact max size is specified in <b>MAX_FILESIZE</b> and may be obtained using the <b>pathconf</b> system call. Any file larger than that cannot be opened on the 32-bit kernel, but can be created and opened on the 64-bit kernel.
<b>EROFS</b>	Named file resides on a read-only file system and write access is required (either the <b>O_WRONLY</b> , <b>O_RDWR</b> , <b>O_CREAT</b> (if the file does not exist), or <b>O_TRUNC</b> flag is set in the <i>OFlag</i> parameter).
<b>ETXTBSY</b>	File is on a physical file system and is already open in a manner (with the <b>O_RSHARE</b> or <b>O_NSHARE</b> flag) that precludes this open; or the <b>O_NSHARE</b> or <b>O_RSHARE</b> flag was requested with the <b>O_NDELAY</b> flag set, and there is a conflicting open on a physical file system.

**Note:** The **E\_OVERFLOW** error code applies to AIX 4.2 and later releases.

**E\_OVERFLOW** A call was made to **open** and **creat** and the file already existed and its size was larger than **OFF\_MAX** and the **O\_LARGEFILE** flag was not set.

The **open**, **openx**, and **creat** subroutines are unsuccessful if one of the following are true:

<b>EFAULT</b>	The <i>Path</i> parameter points outside of the allocated address space of the process.
<b>EINVAL</b>	The value of the <i>OFlag</i> parameter is not valid.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ETXTBSY</b>	The file specified by the <i>Path</i> parameter is a pure procedure (shared text) file that is currently executing, and the <b>O_WRONLY</b> or <b>O_RDWR</b> flag is set in the <i>OFlag</i> parameter.

## Related Information

The **chmod** (“chmod or fchmod Subroutine” on page 148) subroutine, **close** (“close Subroutine” on page 175) subroutine, **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutine, **fcntl**, **dup**, or **dup2** (“fcntl, dup, or dup2 Subroutine” on page 255) subroutine, **fsync** (“fsync or fsync\_range Subroutine” on page 321) subroutine, **ioctl** (“ioctl, ioctlx, ioctl32, or ioctl32x Subroutine” on page 561) subroutine, **lockf** (“lockf, flock, or lockf64 Subroutine” on page 740) subroutine, **lseek** (“lseek, llseek or lseek64 Subroutine” on page 765) subroutine, **read** subroutine, **stat** subroutine, **umask** subroutine, **write** subroutine.

The Input and Output Handling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## **opendir, readdir, telldir, seekdir, rewinddir, closedir, opendir64, readdir64, telldir64, seekdir64, rewinddir64, or closedir64 Subroutine**

### Purpose

Performs operations on directories.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <dirent.h>
```

```
DIR *opendir ( DirectoryName)
const char *DirectoryName;
```

```
struct dirent *readdir ( DirectoryPointer)
DIR *DirectoryPointer;
```

```
long int telldir(DirectoryPointer)
DIR *DirectoryPointer;
```

```
void seekdir(DirectoryPointer,Location)
DIR *DirectoryPointer;
long Location;
```

```
void rewinddir (DirectoryPointer)
DIR *DirectoryPointer;
```

```
int closedir (DirectoryPointer)
DIR *DirectoryPointer;
```

```
DIR *opendir64 ( DirectoryName)
const char *DirectoryName;
```

```
struct dirent64 *readdir64 ( DirectoryPointer)
DIR64 *DirectoryPointer;
```

```
offset_t telldir64(DirectoryPointer)
DIR64 *DirectoryPointer;
```

```
void seekdir64(DirectoryPointer,Location)
DIR64 *DirectoryPointer;
offset_t Location;
```

```
void rewinddir64 (DirectoryPointer)
DIR64 *DirectoryPointer;
```

```
int closedir64 (DirectoryPointer)
DIR64 *DirectoryPointer;
```

## Description

**Attention:** Do not use the **readdir** subroutine in a multithreaded environment. See the multithread alternative in the **readdir\_r** subroutine article.

The **opendir** subroutine opens the directory designated by the *DirectoryName* parameter and associates a directory stream with it.

**Note:** An open directory must always be closed with the **closedir** subroutine to ensure that the next attempt to open that directory is successful.

The **opendir** subroutine also returns a pointer to identify the directory stream in subsequent operations. The null pointer is returned when the directory named by the *DirectoryName* parameter cannot be accessed or when not enough memory is available to hold the entire stream. A successful call to any of the **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) functions closes any directory streams opened in the calling process.

The **readdir** subroutine returns a pointer to the next directory entry. The **readdir** subroutine returns entries for . (dot) and .. (dot dot), if present, but never returns an invalid entry (with *d\_ino* set to 0). When it reaches the end of the directory, or when it detects an invalid **seekdir** operation, the **readdir** subroutine returns the null value. The returned pointer designates data that may be overwritten by another call to the **readdir** subroutine on the same directory stream. A call to the **readdir** subroutine on a different directory stream does not overwrite this data. The **readdir** subroutine marks the *st\_atime* field of the directory for update each time the directory is actually read.

The **telldir** subroutine returns the current location associated with the specified directory stream.

The **seekdir** subroutine sets the position of the next **readdir** subroutine operation on the directory stream. An attempt to seek an invalid location causes the **readdir** subroutine to return the null value the next time it is called. The position should be that returned by a previous **telldir** subroutine call.

The **rewinddir** subroutine resets the position of the specified directory stream to the beginning of the directory.

The **closedir** subroutine closes a directory stream and frees the structure associated with the *DirectoryPointer* parameter. If the **closedir** subroutine is called for a directory that is already closed, the behavior is undefined. To prevent this, always initialize the *DirectoryPointer* parameter to null after closure.

If you use the **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine to create a new process from an existing one, either the parent or the child (but not both) may continue processing the directory stream using the **readdir**, **rewinddir**, or **seekdir** subroutine.

The **opendir64** subroutine is similar to the **opendir** subroutine except that it returns a pointer to an object of type **DIR64**.

**Note:** An open directory by **opendir64** subroutine must always be closed with the **closedir64** subroutine to ensure that the next attempt to open that directory is successful. In addition, it must be operated using the 64-bit interfaces (**readdir64**, **telldir64**, **seekdir64**, **rewinddir64**, and **closedir64**) to obtain the correct directory information.

The **readdir64** subroutine is similar to the **readdir** subroutine except that it returns a pointer to an object of type **struct dirent64**.

The **telldir64** subroutine is similar to the **telldir** subroutine except that it returns the current directory location in an **offset\_t** format.



The **seekdir64** subroutine is similar to the **seekdir** subroutine except that the *Location* parameter is set in the format of **offset\_t**.

The **rewinddir64** subroutine resets the position of the specified directory stream (obtained by the **opendir64** subroutine) to the beginning of the directory.

## Parameters

<i>DirectoryName</i>	Names the directory.
<i>DirectoryPointer</i>	Points to the <b>DIR</b> or <b>DIR64</b> structure of an open directory.
<i>Location</i>	Specifies the offset of an entry relative to the start of the directory.

## Return Values

On successful completion, the **opendir** subroutine returns a pointer to an object of type **DIR**, and the **opendir64** subroutine returns a pointer to an object of type **DIR64**. Otherwise, a null value is returned and the **errno** global variable is set to indicate the error.

On successful completion, the **readdir** subroutine returns a pointer to an object of type **struct dirent**, and the **readdir64** subroutine returns a pointer to an object of type **struct dirent64**. Otherwise, a null value is returned and the **errno** global variable is set to indicate the error. When the end of the directory is encountered, a null value is returned and the **errno** global variable is not changed by this function call.

On successful completion, the **telldir** or **telldir64** subroutine returns the current location associated with the specified directory stream. Otherwise, a null value is returned.

On successful completion, the **closedir** or **closedir64** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

If the **opendir** subroutine is unsuccessful, it returns a null value and sets the **errno** global variable to one of the following values:

<b>EACCES</b>	Indicates that search permission is denied for any component of the <i>DirectoryName</i> parameter, or read permission is denied for the <i>DirectoryName</i> parameter.
<b>ENAMETOOLONG</b>	Indicates that the length of the <i>DirectoryName</i> parameter argument exceeds the <b>PATH_MAX</b> value, or a path-name component is longer than the <b>NAME_MAX</b> value while the <b>POSIX_NO_TRUNC</b> value is in effect.
<b>ENOENT</b>	Indicates that the named directory does not exist.
<b>ENOTDIR</b>	Indicates that a component of the <i>DirectoryName</i> parameter is not a directory.
<b>EMFILE</b>	Indicates that too many file descriptors are currently open for the process.
<b>ENFILE</b>	Indicates that too many file descriptors are currently open in the system.

If the **readdir** or **readdir64** subroutine is unsuccessful, it returns a null value and sets the **errno** global variable to the following value:

**EBADF** Indicates that the *DirectoryPointer* parameter argument does not refer to an open directory stream.

If the **closedir** or **closedir64** subroutine is unsuccessful, it returns a value of -1 and sets the **errno** global variable to the following value:

**EBADF** Indicates that the *DirectoryPointer* parameter argument does not refer to an open directory stream.

## Examples

To search a directory for the entry name:

```
len = strlen(name);
DirectoryPointer = opendir(".");
for (dp = readdir(DirectoryPointer); dp != NULL; dp =
    readdir(DirectoryPointer))
    if (dp->d_namlen == len && !strcmp(dp->d_name, name)) {
        closedir(DirectoryPointer);
        DirectoryPointer=NULL //To prevent multiple closure
        return FOUND;
    }
closedir(DirectoryPointer);
DirectoryPointer=NULL //To prevent multiple closure
```

## Related Information

The **close** (“close Subroutine” on page 175) subroutine, **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutines, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **lseek** (“lseek, llseek or lseek64 Subroutine” on page 765) subroutine, **openx**, **open**, or **creat** (“open, openx, open64, creat, or creat64 Subroutine” on page 935) subroutine, **read**, **readv**, **readx**, or **readvx** subroutine, **scandir** or **alphasort** subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pam\_acct\_mgmt Subroutine

### Purpose

Validates the user’s account.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

int pam_acct_mgmt (PAMHandle, Flags)
pam_handle_t *PAMHandle;
int Flags;
```

### Description

The **pam\_acct\_mgmt** subroutine performs various checks on the user’s account to determine if it is valid. These checks can include account and password expiration, and access restrictions. This subroutine is generally used subsequent to a successful **pam\_authenticate()** call in order to verify whether the authenticated user should be granted access.

### Parameters

**PAMhandle**                      The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.

*Flags*

The *Flags* argument can be a logically OR'd combination of the following:

- **PAM\_SILENT**
  - No messages should be displayed
- **PAM\_DISALLOW\_NULL\_AUTHCHK**
  - Do not authenticate a user with a NULL authentication token.

## Return Values

Upon successful completion, **pam\_acct\_mgmt** returns **PAM\_SUCCESS**. If the routine fails, a different error will be returned, depending on the actual error.

## Error Codes

**PAM\_ACCT\_EXPIRED**

The user's account has expired.

**PAM\_NEW\_AUTHCHK\_REQD**

The user's password needs changed. This is usually due to password aging or because it was last set by an administrator. At this stage most user's can still change their passwords; applications should call **pam\_chauthchk()** and have the user promptly change their password.

**PAM\_AUTHCHK\_EXPIRED**

The user's password has expired. Unlike **PAM\_NEW\_AUTHCHK\_REQD**, the password cannot be changed by the user.

**PAM\_USER\_UNKNOWN**

The user is not known.

**PAM\_OPEN\_ERR**

One of the PAM authentication modules could not be loaded.

**PAM\_SYMBOL\_ERR**

A necessary item is not available to a PAM module.

**PAM\_SERVICE\_ERR**

An error occurred in a PAM module.

**PAM\_SYSTEM\_ERR**

A system error occurred.

**PAM\_BUF\_ERR**

A memory error occurred.

**PAM\_CONV\_ERR**

A conversation error occurred.

**PAM\_PERM\_DENIED**

Access permission was denied to the user.

## Related Information

"pam\_authenticate Subroutine," "pam\_open\_session Subroutine" on page 957, "pam\_setcred Subroutine" on page 961, "pam\_sm\_acct\_mgmt Subroutine" on page 963, "pam\_start Subroutine" on page 970

---

## pam\_authenticate Subroutine

### Purpose

Attempts to authenticate a user through PAM.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
```

```
int pam_authenticate (PAMHandle, Flags)
pam_handle_t *PAMHandle;
int Flags;
```

## Description

The **pam\_authenticate** subroutine authenticates a user through PAM. The authentication method used is determined by the authentication modules configured in the **/etc/pam.conf** stack. Most authentication requires a password or other user input but is dependent on the modules in use.

Before attempting authentication through **pam\_authenticate**, ensure that all of the applicable PAM information has been set through the initial call to **pam\_start()** and subsequent calls to **pam\_set\_item()**. If any necessary information is not set, PAM modules can prompt the user for information through the routine defined in **PAM\_CONV**. If required information is not provided and **PAM\_CONV** is not set, the authentication fails.

On failure, it is the responsibility of the calling application to maintain a count of authentication attempts and to reinvoke the subroutine if the count has not exceeded a defined limit. Some authentication modules maintain an internal count and return **PAM\_MAXTRIES** if the limit is reached. After the stack of authentication modules has finished with either success or failure, **PAM\_AUTHTOK** is cleared in the handle.

## Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>Flags</i>	The Flags argument can be a logically OR'd combination of the following: <ul style="list-style-type: none"><li>• <b>PAM_SILENT</b><ul style="list-style-type: none"><li>– No messages should be displayed</li></ul></li><li>• <b>PAM_DISALLOW_NULL_AUTHTOK</b><ul style="list-style-type: none"><li>– Do not authenticate a user with a NULL authentication token.</li></ul></li></ul>

## Return Values

Upon successful completion, **pam\_authenticate** returns **PAM\_SUCCESS**. If the routine fails, a different error will be returned, depending on the actual error.

## Error Codes

<b>PAM_AUTH_ERR</b>	An error occurred in authentication, usually because of an invalid authentication token.
<b>PAM_CRED_INSUFFICIENT</b>	The user has insufficient credentials to access the authentication data.
<b>PAM_AUTHINFO_UNAVAIL</b>	The authentication information cannot be retrieved.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_MAXTRIES</b>	The maximum number of authentication retries has been reached.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_acct\_mgmt Subroutine” on page 946, “pam\_get\_user Subroutine” on page 954, “pam\_open\_session Subroutine” on page 957, “pam\_set\_item Subroutine” on page 960, “pam\_setcred Subroutine” on page 961, “pam\_sm\_authenticate Subroutine” on page 964, “pam\_start Subroutine” on page 970

---

## pam\_chauthtok Subroutine

### Purpose

Changes the user’s authentication token (typically passwords).

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

int pam_chauthtok (PAMHandle, Flags)
pam_handle_t *PAMHandle;
int Flags;
```

### Description

The **pam\_chauthtok** subroutine changes a user’s authentication token through the PAM framework. Prior to changing the password, the subroutine performs preliminary tests to ensure that necessary hosts and information, depending on the password service, are there. If any of these tests fail, **PAM\_TRY\_AGAIN** is returned. To request information from the user, **pam\_chauthtok** can use the conversation function that is defined in the PAM handle, *PAMHandle*. After the subroutine is finished, the values of **PAM\_AUTH Tok** and **PAM\_OLDAUTH Tok** are cleared in the handle for added security.

### Parameters

*PAMhandle*

The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.

*Flags*

The Flags argument can be a logically OR’d combination of the following:

- **PAM\_SILENT**
  - No messages should be displayed
- **PAM\_CHANGE\_EXPIRED\_AUTH Tok**
  - Only expired passwords should be changed. If this flag is not included, all users using the related password service are forced to update their passwords. This is typically used by a login application after determining password expiration. It should not generally be used by applications dedicated to changing passwords.

### Return Values

Upon successful completion, **pam\_chauthtok** returns **PAM\_SUCCESS** and the authentication token of the user, as defined for a given password service, is changed. If the routine fails, a different error is returned, depending on the actual error.

### Error Codes

**PAM\_AUTH Tok\_ERR**

A failure occurred while updating the authentication token.

<b>PAM_TRY_AGAIN</b>	Preliminary checks for changing the password have failed. Try again later.
<b>PAM_AUTHTOK_RECOVERY_ERR</b>	An error occurred while trying to recover the authentication information.
<b>PAM_AUTHTOK_LOCK_BUSY</b>	Cannot get the authentication token lock. Try again later.
<b>PAM_AUTHTOK_DISABLE_AGING</b>	Authentication token aging checks are disabled and were not performed.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_acct\_mgmt Subroutine” on page 946, “pam\_authenticate Subroutine” on page 947, “pam\_open\_session Subroutine” on page 957, “pam\_setcred Subroutine” on page 961, “pam\_sm\_chauthtok Subroutine” on page 965, “pam\_start Subroutine” on page 970

---

## pam\_close\_session Subroutine

### Purpose

Ends a currently open PAM user session.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

int pam_close_session (PAMHandle, Flags)
pam_handle_t *PAMHandle;
int Flags;
```

### Description

The **pam\_close\_session** subroutine ends a PAM user session started by **pam\_open\_session()**.

### Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>Flags</i>	The following flag may be set: <ul style="list-style-type: none"> <li>• <b>PAM_SILENT</b> <ul style="list-style-type: none"> <li>– No messages should be displayed</li> </ul> </li> </ul>

## Return Values

Upon successful completion, `pam_close_session` returns `PAM_SUCCESS`. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

<code>PAM_SESSION_ERR</code>	An error occurred while creating/removing an entry for the new session.
<code>PAM_USER_UNKNOWN</code>	The user is not known.
<code>PAM_OPEN_ERR</code>	One of the PAM authentication modules could not be loaded.
<code>PAM_SYMBOL_ERR</code>	A necessary item is not available to a PAM module.
<code>PAM_SERVICE_ERR</code>	An error occurred in a PAM module.
<code>PAM_SYSTEM_ERR</code>	A system error occurred.
<code>PAM_BUF_ERR</code>	A memory error occurred.
<code>PAM_CONV_ERR</code>	A conversation error occurred.
<code>PAM_PERM_DENIED</code>	Access permission was denied to the user.

## Related Information

“`pam_open_session` Subroutine” on page 957, “`pam_sm_close_session` Subroutine” on page 967, “`pam_start` Subroutine” on page 970

---

## `pam_end` Subroutine

### Purpose

Ends an existing PAM authentication session.

### Library

PAM Library (`libpam.a`)

### Syntax

```
#include <security/pam_appl.h>

int pam_end (PAMHandle, Status)
pam_handle_t *PAMHandle;
int Status;
```

### Description

The `pam_end` subroutine finishes and cleans up the authentication session represented by the PAM handle `PAMHandle`. `Status` denotes the current state of the `PAMHandle` and is passed through to a `cleanup()` function so that the memory used during that session can be properly unallocated. The `cleanup()` function can be set in the `PAMHandle` by PAM modules through the `pam_set_data()` routine. Upon completion of the subroutine, the PAM handle and associated memory is no longer valid.

### Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <code>pam_start()</code> .
<i>Status</i>	The state of the last PAM call. Some modules need to be cleaned according to error codes.

## Return Values

Upon successful completion, **pam\_end** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.

## Related Information

“pam\_start Subroutine” on page 970

---

## pam\_get\_data Subroutine

### Purpose

Retrieves information for a specific PAM module for this PAM session.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

int pam_get_data (PAMHandle, ModuleDataName, Data)
pam_handle_t *PAMHandle;
const char *ModuleDataName;
void **Data;
```

### Description

The **pam\_get\_data** subroutine is used to retrieve module-specific data from the PAM handle. This subroutine is used by modules and should not be called by applications. If the *ModuleDataName* identifier exists, the reference for its data is returned in *Data*. If the identifier does not exist, a NULL reference is returned in *Data*. The caller should not modify or free the memory returned in *Data*. Instead, a cleanup function should be specified through a call to **pam\_set\_data()**. The cleanup function will be called when **pam\_end()** is invoked in order to free any memory allocated.

### Parameters

<i>PAMHandle</i> (in)	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>ModuleDataName</i>	A unique identifier for <i>Data</i> .
<i>Data</i>	Returned reference to the data denoted by <i>ModuleDataName</i> .

### Return Values

Upon successful completion, **pam\_get\_data** returns **PAM\_SUCCESS**. If *ModuleDataName* exists and **pam\_get\_data** completes successfully, *Data* will be a valid reference. Otherwise, *Data* will be NULL. If the routine fails, either **PAM\_SYSTEM\_ERR**, **PAM\_BUF\_ERR**, or **PAM\_NO\_MODULE\_DATA** is returned, depending on the actual error.



## Error Codes

<code>PAM_SYSTEM_ERR</code>	A system error occurred.
<code>PAM_BUF_ERR</code>	A memory error occurred.
<code>PAM_NO_MODULE_DATA</code>	No module-specific data was found.

## Related Information

“`pam_get_item` Subroutine,” “`pam_getenv` Subroutine” on page 955, “`pam_getenvlist` Subroutine” on page 956, “`pam_set_data` Subroutine” on page 959

---

## `pam_get_item` Subroutine

### Purpose

Retrieves an item or information for this PAM session.

### Library

PAM Library (`libpam.a`)

### Syntax

```
#include <security/pam_appl.h>

int pam_get_item (PAMHandle, ItemType, Item)
pam_handle_t *PAMHandle;
int ItemType;
void **Item;
```

### Description

The `pam_get_item` subroutine returns the item requested by the *ItemType*. Any items returned by `pam_get_item` should not be modified or freed. They can be later used by PAM and will be cleaned-up by `pam_end()`. If a requested *ItemType* is not found, a NULL reference will be returned in *Item*.

### Parameters

*PAMhandle*                      The PAM handle representing the current user authentication session. This handle is obtained by a call to `pam_start()`.

*ItemType*

The type of item that is being requested. The following values are valid item types:

- **PAM\_SERVICE**
  - The service name requesting this PAM session.
- **PAM\_USER**
  - The user name of the user being authenticated.
- **PAM\_AUTHTOK**
  - The user's current authentication token (password).
- **PAM\_OLDAUTHOK**
  - The user's old authentication token (old password).
- **PAM\_TTY**
  - The terminal name.
- **PAM\_RHOST**
  - The name of the remote host.
- **PAM\_RUSER**
  - The name of the remote user.
- **PAM\_CONV**
  - The **pam\_conv** structure for conversing with the user.
- **PAM\_USER\_PROMPT**
  - The default prompt for the user (used by **pam\_get\_user()**).

For security, **PAM\_AUTHTOK** and **PAM\_OLDAUTHOK** are only available to PAM modules.

*Item*

The return value, holding a reference to a pointer of the requested *ItemType*.

## Return Values

Upon successful completion, **pam\_get\_item** returns **PAM\_SUCCESS**. Also, the address of a reference to the requested object is returned in *Item*. If the requested item was not found, a NULL reference is returned. If the routine fails, either **PAM\_SYSTEM\_ERR** or **PAM\_BUF\_ERR** is returned and *Item* is set to a NULL pointer.

## Error Codes

<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_SYMBOL_ERR</b>	Symbol not found.

## Related Information

“pam\_get\_data Subroutine” on page 952, “pam\_getenv Subroutine” on page 955, “pam\_get\_user Subroutine,” “pam\_getenvlist Subroutine” on page 956, “pam\_set\_item Subroutine” on page 960

---

## pam\_get\_user Subroutine

### Purpose

Gets the user's name from the PAM handle or through prompting for input.

### Library

PAM Library (**libpam.a**)

## Syntax

```
#include <security/pam_appl.h>

int pam_get_user (PAMHandle, User, PromptMsg)
pam_handle_t *PAMHandle;
int User;
void **PromptMsg;
```

## Description

The **pam\_get\_user** subroutine returns the user name currently stored in the PAM handle, *PAMHandle*. If the user name has not already been set through **pam\_start()** or **pam\_set\_item()**, the subroutine displays the string specified by *PromptMsg*, to prompt for the user name through the conversation function. If *PromptMsg* is NULL, the value of **PAM\_USER\_PROMPT** set through a call to **pam\_set\_item()** is used. If both *PromptMsg* and **PAM\_USER\_PROMPT** are NULL, PAM defaults to use the following string:

Please enter user name:

After the user name has been retrieved, it is set in the PAM handle and is also returned to the caller in the *User* argument. The caller should not change or free *User*, as cleanup will be handled by **pam\_end()**.

## Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>User</i>	The user name retrieved from the PAM handle or provided by the user.
<i>PromptMsg</i>	The prompt to be displayed if a user name is required and has not been already set.

## Return Values

Upon successful completion, **pam\_get\_user** returns **PAM\_SUCCESS**. Also, a reference to the user name is returned in *User*. If the routine fails, either **PAM\_SYSTEM\_ERR**, **PAM\_BUF\_ERR**, or **PAM\_CONV\_ERR** is returned, depending on what the actual error was, and a NULL reference in *User* is returned.

## Error Codes

<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error or failure.

## Related Information

“pam\_end Subroutine” on page 951, “pam\_get\_item Subroutine” on page 953, “pam\_set\_item Subroutine” on page 960

---

## pam\_getenv Subroutine

### Purpose

Returns the value of a defined PAM environment variable.

### Library

PAM Library (**libpam.a**)

## Syntax

```
#include <security/pam_appl.h>

char *pam_getenv (PAMHandle, VarName)
pam_handle_t *PAMHandle;
char *VarName;
```

## Description

The **pam\_getenv** subroutine retrieves the value of the PAM environment variable *VarName* stored in the PAM handle *PAMHandle*. Environment variables can be defined through the **pam\_putenv()** call. If *VarName* is defined, its value is returned in memory allocated by the library; it is the caller's responsibility to free this memory. Otherwise, a NULL pointer is returned.

## Parameters

<i>PAMHandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>VarName</i>	The name of the PAM environment variable to get the value for.

## Return Values

Upon successful completion, **pam\_getenv** returns the value of the *VarName* PAM environment variable. If the routine fails or *VarName* is not defined, NULL is returned.

## Related Information

“pam\_getenvlist Subroutine,” “pam\_putenv Subroutine” on page 958

---

## pam\_getenvlist Subroutine

### Purpose

Returns a list of all of the defined PAM environment variables and their values.

### Library

PAM Library (**libpam.a**)

## Syntax

```
#include <security/pam_appl.h>

char **pam_getenvlist (PAMHandle)
pam_handle_t *PAMHandle;
```

## Description

The **pam\_getenvlist** subroutine returns a pointer to a list of the currently defined environment variables in the PAM handle, *PAMHandle*. Environment variables can be set through calls to the **pam\_putenv()** subroutine. The library returns the environment in an allocated array in which the last entry of the array is NULL. The caller is responsible for freeing the memory of the returned list.

## Parameters

<i>PAMHandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
------------------	--

## Return Values

Upon successful completion, **pam\_getenvlist** returns a pointer to a list of strings, one for each currently defined PAM environment variable. Each string is of the form *VARIABLE=VALUE*, where *VARIABLE* is the name of the variable and *VALUE* is its value. This list is terminated with a NULL entry. If the routine fails or there are no PAM environment variables defined, a NULL reference is returned. The caller is responsible for freeing the memory of the returned value.

## Related Information

“pam\_getenv Subroutine” on page 955, “pam\_putenv Subroutine” on page 958

---

## pam\_open\_session Subroutine

### Purpose

Opens a new PAM user session.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

int pam_open_session (PAMHandle, Flags)
pam_handle_t *PAMHandle;
int Flags;
```

### Description

The **pam\_open\_session** subroutine opens a new user session for an authenticated PAM user. A call to **pam\_authenticate()** is typically made prior to invoking this subroutine. Applications that open a user session should subsequently close the session with **pam\_close\_session()** when the session has ended.

### Parameters

*PAMhandle*                      The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.

*Flags*                              The flags are used to set pam\_acct\_mgmt options. The recognized flags are:

- **PAM\_SILENT**
  - No messages should be displayed

### Return Values

Upon successful completion, **pam\_open\_session** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

### Error Codes

<b>PAM_SESSION_ERR</b>	An error occurred while creating/removing an entry for the new session.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.

<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_authenticate Subroutine” on page 947, “pam\_close\_session Subroutine” on page 950, “pam\_sm\_open\_session Subroutine” on page 968, “pam\_start Subroutine” on page 970

---

## pam\_putenv Subroutine

### Purpose

Defines a PAM environment variable.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

int pam_putenv (PAMHandle, NameValue)
pam_handle_t *PAMHandle;
const char *NameValue;
```

### Description

The **pam\_putenv** subroutine sets and deletes environment variables in the PAM handle, *PAMHandle*. Applications can retrieve the defined variables by calling **pam\_getenv()** or **pam\_getenvlist()** and add them to the user’s session. If a variable with the same name is already defined, the old value is replaced by the new value.

### Parameters

<i>PAMHandle</i>	The PAM authentication handle, obtained from a previous call to <b>pam_start()</b> .
<i>NameValue</i>	A string of the form <i>name=value</i> to be stored in the environment section of the PAM handle. The following behavior is exhibited with regards to the format of the passed-in string: <ul style="list-style-type: none"> <li><i>NAME=VALUE</i> Creates or overwrites the value for the variable in the environment.</li> <li><i>NAME=</i> Sets the variable to the empty string.</li> <li><i>NAME</i> Deletes the variable from the environment, if it is currently defined.</li> </ul>

### Return Values

Upon successful completion, **pam\_putenv** returns **PAM\_SUCCESS**. If the routine fails, either **PAM\_SYSTEM\_ERR** or **PAM\_BUF\_ERR** is returned, depending on the actual error.

## Error Codes

<code>PAM_SYSTEM_ERR</code>	A system error occurred.
<code>PAM_BUF_ERR</code>	A memory error occurred.

## Related Information

“`pam_getenv` Subroutine” on page 955, “`pam_getenvlist` Subroutine” on page 956, “`pam_start` Subroutine” on page 970

---

## `pam_set_data` Subroutine

### Purpose

Sets information for a specific PAM module for the active PAM session.

### Library

PAM Library (`libpam.a`)

### Syntax

```
#include <security/pam_appl.h>

int pam_set_data (PAMHandle, ModuleDataName, Data, *(cleanup)(pam_handle_t *pamh, void *data,
                    int pam_end_status))
pam_handle_t *PAMHandle;
const char *ModuleDataName;
void *Data;
void *(cleanup)(pam_handle_t *pamh, void *data, int pam_end_status);
```

### Description

The `pam_set_data` subroutine allows for the setting and updating of module-specific data within the PAM handle, `PAMHandle`. The `ModuleDataName` argument serves to uniquely identify the data, `Data`. Stored information can be retrieved by specifying `ModuleDataName` and passing it, along with the appropriate PAM handle, to `pam_get_data()`. The `cleanup` argument is a pointer to a function that is called to free allocated memory used by the `Data` when `pam_end()` is invoked. If data is already associated with `ModuleDataName`, PAM does a cleanup of the old data, overwrites it with `Data`, and replaces the old `cleanup` function. If the information being set is of a known PAM item type, use the `pam_putenv` subroutine instead.

### Parameters

<code>PAMHandle</code>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <code>pam_start()</code> .
<code>ModuleDataName</code>	A unique identifier for <code>Data</code> .
<code>Data</code>	A reference to the data denoted by <code>ModuleDataName</code> .
<code>cleanup</code>	A function pointer that is called by <code>pam_end()</code> to clean up all allocated memory used by <code>Data</code> .

### Return Values

Upon successful completion, `pam_set_data` returns `PAM_SUCCESS`. If the routine fails, either `PAM_SYSTEM_ERR` or `PAM_BUF_ERR` is returned, depending on the actual error.

## Error Codes

<code>PAM_SYSTEM_ERR</code>	A system error occurred.
<code>PAM_BUF_ERR</code>	A memory error occurred.

## Related Information

“`pam_end` Subroutine” on page 951, “`pam_get_data` Subroutine” on page 952, “`pam_get_item` Subroutine” on page 953, “`pam_set_item` Subroutine”

---

## `pam_set_item` Subroutine

### Purpose

Sets the value of an item for this PAM session.

### Library

PAM Library (`libpam.a`)

### Syntax

```
#include <security/pam_appl.h>

int pam_set_item (PAMHandle, ItemType, Item)
pam_handle_t *PAMHandle;
int ItemType;
void **Item;
```

### Description

The `pam_set_item` subroutine allows for the setting and updating of a set of known PAM items. The item value is stored within the PAM handle, *PAMHandle*. If a previous value exists for the item type, *ItemType*, then the old value is overwritten with the new value, *Item*.

### Parameters

*PAMhandle*                    The PAM handle representing the current user authentication session. This handle is obtained by a call to `pam_start()`.



*ItemType*

The type of item that is being requested. The following values are valid item types:

- **PAM\_SERVICE**
  - The service name requesting this PAM session.
- **PAM\_USER**
  - The user name of the user being authenticated.
- **PAM\_AUTHTOK**
  - The user's current authentication token. Interpreted as the new authentication token by password modules.
- **PAM\_OLDAUTHOK**
  - The user's old authentication token. Interpreted as the current authentication token by password modules.
- **PAM\_TTY**
  - The terminal name.
- **PAM\_RHOST**
  - The name of the remote host.
- **PAM\_RUSER**
  - The name of the remote user.
- **PAM\_CONV**
  - The **pam\_conv** structure for conversing with the user.
- **PAM\_USER\_PROMPT**
  - The default prompt for the user (used by **pam\_get\_user()**).

For security, **PAM\_AUTHTOK** and **PAM\_OLDAUTHOK** are only available to PAM modules.

*Item*

The value that the *ItemType* is set to.

## Return Values

Upon successful completion, **pam\_set\_item** returns **PAM\_SUCCESS**. If the routine fails, either **PAM\_SYSTEM\_ERR** or **PAM\_BUF\_ERR** is returned, depending on what the actual error was.

## Error Codes

<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_SYMBOL_ERR</b>	Symbol not found.

## Related Information

“pam\_get\_item Subroutine” on page 953, “pam\_get\_user Subroutine” on page 954

---

## pam\_setcred Subroutine

### Purpose

Establishes, changes, or removes user credentials for authentication.

### Library

PAM Library (**libpam.a**)

## Syntax

```
#include <security/pam_appl.h>

int pam_setcred (PAMHandle, Flags)
pam_handle_t *PAMHandle;
int Flags;
```

## Description

The **pam\_setcred** subroutine allows for the credentials of the PAM user for the current PAM session to be modified. Functions such as establishing, deleting, renewing, and refreshing credentials are defined.

## Parameters

*PAMhandle*

The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.

*Flags*

The flags are used to set **pam\_setcred** options. The recognized flags are:

- **PAM\_SILENT**
  - No messages should be displayed.
- **PAM\_ESTABLISH\_CRED\***
  - Sets the user's credentials. This is the default.
- **PAM\_DELETE\_CRED\***
  - Removes the user credentials.
- **PAM\_REINITIALIZE\_CRED\***
  - Renews the user credentials.
- **PAM\_REFRESH\_CRED\***
  - Refresh the user credentials, extending their lifetime.

\*Mutually exclusive but may be logically OR'd with **PAM\_SILENT**. If one of them is not set, **PAM\_ESTABLISH\_CRED** is assumed.

## Return Values

Upon successful completion, **pam\_setcred** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

<b>PAM_CRED_UNAVAIL</b>	The user credentials cannot be found.
<b>PAM_CRED_EXPIRED</b>	The user's credentials have expired.
<b>PAM_CRED_ERR</b>	A failure occurred while setting user credentials.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_acct\_mgmt Subroutine” on page 946, “pam\_authenticate Subroutine” on page 947, “pam\_open\_session Subroutine” on page 957, “pam\_sm\_setcred Subroutine” on page 969, “pam\_start Subroutine” on page 970

---

## pam\_sm\_acct\_mgmt Subroutine

### Purpose

PAM module implementation for **pam\_acct\_mgmt()**.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
#include <security/pam_modules.h>

int pam_sm_acct_mgmt (PAMHandle, Flags, Argc, Argv)
pam_handle_t *PAMHandle;
int Flags;
int Argc;
const char **Argv;
```

### Description

The **pam\_sm\_acct\_mgmt** subroutine is invoked by the PAM library in response to a call to **pam\_acct\_mgmt**. The **pam\_sm\_acct\_mgmt** subroutine performs the account and password validation for a user and is associated with the “account” service in the PAM configuration file. It is up to the module writers to implement their own service-dependent version of **pam\_sm\_acct\_mgmt**, if the module requires this feature. Actual checks performed are at the discretion of the module writer but typically include checks such as password expiration and login time validation.

### Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>Flags</i>	The <i>Flags</i> argument can be a logically OR’d combination of the following: <ul style="list-style-type: none"><li>• <b>PAM_SILENT</b><ul style="list-style-type: none"><li>– No messages should be displayed.</li></ul></li><li>• <b>PAM_DISALLOW_NULL_AUTHCHK</b><ul style="list-style-type: none"><li>– Do not authenticate a user with a NULL authentication token.</li></ul></li></ul>
<i>Argc</i>	The number of module options specified in the PAM configuration file.
<i>Argv</i>	The module options specified in the PAM configuration file. These options are module-dependent. Any modules receiving invalid options should ignore them.

### Return Values

Upon successful completion, **pam\_sm\_acct\_mgmt** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

### Error Codes

<b>PAM_ACCT_EXPIRED</b>	The user’s account has expired.
-------------------------	---------------------------------

**PAM\_NEW\_AUTHTOKEN\_REQD**

The user's password needs to be changed. This is usually due to password aging or because it was last set by the system administrator. At this stage, most users can still change their passwords. Applications should call **pam\_chauthtok()** and have the users change their password.

**PAM\_AUTHTOK\_EXPIRED**

The user's password has expired. Unlike **PAM\_NEW\_AUTHTOKEN\_REQD**, the password cannot be changed by the user.

**PAM\_USER\_UNKNOWN**

The user is not known.

**PAM\_OPEN\_ERR**

One of the PAM authentication modules could not be loaded.

**PAM\_SYMBOL\_ERR**

A necessary item is not available to a PAM module.

**PAM\_SERVICE\_ERR**

An error occurred in a PAM module.

**PAM\_SYSTEM\_ERR**

A system error occurred.

**PAM\_BUF\_ERR**

A memory error occurred.

**PAM\_CONV\_ERR**

A conversation error occurred.

**PAM\_PERM\_DENIED**

Access permission was denied to the user.

## Related Information

“pam\_acct\_mgmt Subroutine” on page 946, “pam\_authenticate Subroutine” on page 947, “pam\_start Subroutine” on page 970

---

## pam\_sm\_authenticate Subroutine

### Purpose

PAM module-specific implementation of **pam\_authenticate()**.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
#include <security/pam_modules.h>

int pam_sm_authenticate (PAMHandle, Flags, Argc, Argv)
pam_handle_t *PAMHandle;
int Flags;
int Argc;
const char **Argv;
```

### Description

When an application invokes **pam\_authenticate()**, the PAM Framework calls **pam\_sm\_authenticate** for each module in the authentication module stack. This allows all the PAM module authors to implement their own authenticate routine. **pam\_authenticate** and **pam\_sm\_authenticate** provide an authentication service to verify that the user is allowed access.

### Parameters

*PAMhandle*                    The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.

*Flags*

The flags are used to set **pam\_acct\_mgmt** options. The recognized flags are:

- **PAM\_SILENT**
  - No messages should be displayed.
- **PAM\_DISALLOW\_NULL\_AUTHCHK**
  - Do not authenticate a user with a NULL authentication token.

*Argc*

The number of module options defined.

*Argv*

The module options. These options are module-dependent. Any modules receiving invalid options should ignore them.

## Return Values

Upon successful completion, **pam\_sm\_authenticate** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

**PAM\_AUTH\_ERR**

An error occurred in authentication, usually because of an invalid authentication token.

**PAM\_CRED\_INSUFFICIENT**

The user has insufficient credentials to access the authentication data.

**PAM\_AUTHINFO\_UNAVAIL**

The authentication information cannot be retrieved.

**PAM\_USER\_UNKNOWN**

The user is not known.

**PAM\_MAXTRIES**

The maximum number of authentication retries has been reached.

**PAM\_OPEN\_ERR**

One of the PAM authentication modules could not be loaded.

**PAM\_SYMBOL\_ERR**

A necessary item is not available to a PAM module.

**PAM\_SERVICE\_ERR**

An error occurred in a PAM module.

**PAM\_SYSTEM\_ERR**

A system error occurred.

**PAM\_BUF\_ERR**

A memory error occurred.

**PAM\_CONV\_ERR**

A conversation error occurred.

**PAM\_PERM\_DENIED**

Access permission was denied to the user.

## Related Information

“pam\_authenticate Subroutine” on page 947

---

## pam\_sm\_chauthtok Subroutine

### Purpose

PAM module-specific implementation of **pam\_chauthtok()**.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
#include <security/pam_modules.h>
```

```
int pam_sm_chauthtok (PAMHandle, Flags, Argc, Argv)
pam_handle_t *PAMHandle;
```

```
int Flags;
int Argc;
const char **Argv;
```

## Description

When an application invokes **pam\_chauthtok()**, the PAM Framework calls **pam\_sm\_chauthtok** for each module in the password module stack. The **pam\_sm\_chauthtok** module interface is intended to change the user's password or authentication token. Before any password is changed, **pam\_sm\_chauthtok** performs preliminary tests to ensure necessary hosts and information, depending on the password service, are there. If **PAM\_PRELIM\_CHECK** is specified, only these preliminary checks are done. If successful, the authentication token is ready to be changed. If the **PAM\_UPDATE\_AUTHHTOK** flag is passed in, **pam\_sm\_chauthtok** should take the next step and change the user's authentication token. If the **PAM\_CHANGE\_EXPIRED\_AUTHHTOK** flag is set, the module should check the authentication token for aging and expiration. If the user's authentication token is aged or expired, the module should store that information by passing it to **pam\_set\_data()**. Otherwise, the module should exit and return **PAM\_IGNORE**. Required information is obtained through the PAM handle or by prompting the user by way of **PAM\_CONV**.

## Parameters

*PAMhandle*

The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.

*Flags*

The flags are used to set **pam\_acct\_mgmt** options. The recognized flags are:

- **PAM\_SILENT**
  - No messages should be displayed.
- **PAM\_CHANGE\_EXPIRED\_AUTHHTOK**
  - Only expired passwords should be changed. If this flag is not included, all users using the related password service are forced to update their passwords.
- **PAM\_PRELIM\_CHECK\***
  - Only perform preliminary checks to see if the password can be changed, but do not change it.
- **PAM\_UPDATE\_AUTHHTOK\***
  - Perform all necessary checks, and if possible, change the user's password/ authentication token.

\* **PAM\_PRELIM\_CHECK** and **PAM\_UPDATE\_AUTHHTOK** are mutually exclusive.

*Argc*

The number of module options defined.

*Argv*

The module options. These options are module-dependent. Any modules receiving invalid options should ignore them.

## Return Values

Upon successful completion, **pam\_sm\_chauthtok** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

**PAM\_AUTHHTOK\_ERR**

A failure occurred while updating the authentication token.

**PAM\_TRY\_AGAIN**

Preliminary checks for changing the password have failed. Try again later.

**PAM\_AUTHHTOK\_RECOVERY\_ERR**

An error occurred while trying to recover the authentication information.

<b>PAM_AUTHTOK_LOCK_BUSY</b>	Cannot get the authentication token lock. Try again later
<b>PAM_AUTHTOK_DISABLE_AGING</b>	Authentication token aging checks are disabled and were not performed.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_chauthtok Subroutine” on page 949

---

## pam\_sm\_close\_session Subroutine

### Purpose

PAM module-specific implementation to close a session previously opened by **pam\_sm\_open\_session()**.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
#include <security/pam_modules.h>

int pam_sm_close_session (PAMHandle, Flags, Argc, Argv)
pam_handle_t *PAMHandle;
int Flags;
int Argc;
const char **Argv;
```

### Description

When an application invokes **pam\_close\_session()**, the PAM Framework calls **pam\_sm\_close\_session** for each module in the session module stack. The **pam\_sm\_close\_session** module interface is intended to clean up and terminate any user session started by **pam\_sm\_open\_session()**.

### Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>Flags</i>	The flags are used to set <b>pam_acct_mgmt</b> options. The recognized flag is: <ul style="list-style-type: none"> <li>• <b>PAM_SILENT</b> <ul style="list-style-type: none"> <li>– No messages should be displayed.</li> </ul> </li> </ul>
<i>Argc</i>	The number of module options defined.
<i>Argv</i>	The module options. These options are module-dependent. Any modules receiving invalid options should ignore them.

## Return Values

Upon successful completion, **pam\_sm\_close\_session** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

<b>PAM_SESSION_ERR</b>	An error occurred while creating or removing an entry for the new session.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_close\_session Subroutine” on page 950, “pam\_sm\_open\_session Subroutine”

---

## pam\_sm\_open\_session Subroutine

### Purpose

PAM module-specific implementation of **pam\_open\_session**.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
#include <security/pam_modules.h>

int pam_sm_open_session (PAMHandle, Flags, Argc, Argv)
pam_handle_t *PAMHandle;
int Flags;
int Argc;
const char **Argv;
```

### Description

When an application invokes **pam\_open\_session()**, the PAM Framework calls **pam\_sm\_open\_session** for each module in the session module stack. The **pam\_sm\_open\_session** module interface starts a new user session for an authenticated PAM user. All session-specific information and memory used by opening a session should be cleaned up by **pam\_sm\_close\_session()**.

### Parameters

*PAMhandle*                      The PAM handle representing the current user authentication session. This handle is obtained by a call to **pam\_start()**.



*Flags* The flags are used to set **pam\_acct\_mgmt** options. The recognized flag is:

- **PAM\_SILENT**
  - No messages should be displayed.

*Argc* The number of module options defined.

*Argv* The module options. These options are module-dependent. Any modules receiving invalid options should ignore them.

## Return Values

Upon successful completion, **pam\_sm\_open\_session** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

<b>PAM_SESSION_ERR</b>	An error occurred while creating or removing an entry for the new session.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

“pam\_open\_session Subroutine” on page 957, “pam\_sm\_close\_session Subroutine” on page 967

---

## pam\_sm\_setcred Subroutine

### Purpose

PAM module-specific implementation of **pam\_setcred**.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>
#include <security/pam_modules.h>

int pam_sm_setcred (PAMHandle, Flags, Argc, Argv)
pam_handle_t *PAMHandle;
int Flags;
int Argc;
const char **Argv;
```

### Description

When an application invokes **pam\_setcred()**, the PAM Framework calls **pam\_sm\_setcred** for each module in the authentication module stack. The **pam\_sm\_setcred** module interface allows for the setting of module-specific credentials in the PAM handle. The user’s credentials should be set based upon the user’s authentication state. This information can usually be retrieved with a call to **pam\_get\_data()**.

## Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>Flags</i>	The flags are used to set <b>pam_setcred</b> options. The recognized flags are: <ul style="list-style-type: none"><li>• <b>PAM_SILENT</b><ul style="list-style-type: none"><li>– No messages should be displayed.</li></ul></li><li>• <b>PAM_ESTABLISH_CRED*</b><ul style="list-style-type: none"><li>– Sets the user's credentials. This is the default.</li></ul></li><li>• <b>PAM_DELETE_CRED*</b><ul style="list-style-type: none"><li>– Removes the user credentials.</li></ul></li><li>• <b>PAM_REINITIALIZE_CRED*</b><ul style="list-style-type: none"><li>– Renews the user credentials.</li></ul></li><li>• <b>PAM_REFRESH_CRED*</b><ul style="list-style-type: none"><li>– Refreshes the user credentials, extending their lifetime.</li></ul></li></ul> <p>*Mutually exclusive. If one of them is not set, <b>PAM_ESTABLISH_CRED</b> is assumed.</p>
<i>Argc</i>	The number of module options defined.
<i>Argv</i>	The module options. These options are module-dependent. Any modules receiving invalid options should ignore them.

## Return Values

Upon successful completion, **pam\_sm\_setcred** returns **PAM\_SUCCESS**. If the routine fails, a different error is returned, depending on the actual error.

## Error Codes

<b>PAM_CRED_UNAVAIL</b>	The user credentials cannot be found.
<b>PAM_CRED_EXPIRED</b>	The user's credentials have expired.
<b>PAM_CRED_ERR</b>	A failure occurred while setting user credentials.
<b>PAM_USER_UNKNOWN</b>	The user is not known.
<b>PAM_OPEN_ERR</b>	One of the PAM authentication modules could not be loaded.
<b>PAM_SYMBOL_ERR</b>	A necessary item is not available to a PAM module.
<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.
<b>PAM_CONV_ERR</b>	A conversation error occurred.
<b>PAM_PERM_DENIED</b>	Access permission was denied to the user.

## Related Information

"pam\_setcred Subroutine" on page 961

---

## pam\_start Subroutine

### Purpose

Initiates a new PAM user authentication session.

### Library

PAM Library (**libpam.a**)

## Syntax

```
#include <security/pam_appl.h>

int pam_start (Service, User, Conversation, PAMHandle)
const char *Service;
const char *User;
const struct pam_conv *Conversation;
pam_handle_t **PAMHandle;
```

## Description

The **pam\_start** subroutine begins a new PAM session for authentication within one of the four realms of the PAM environment [authentication, account, session, password]. This routine is called only at the start of the session, not at the start of each module comprising the session. The PAM handle, *PAMHandle*, returned by this subroutine is subsequently used by other PAM routines. The handle must be cleaned up at the end of use, which can easily be done by passing it as an argument to **pam\_end**.

## Parameters

<i>Service</i>	The name of the service initiating this PAM session.
<i>User</i>	The user who is being authenticated.

## Conversation

The PAM conversation struct enabling communication with the user. This structure, **pam\_conv**, consists of a pointer to a conversation function, as well as a pointer to application data.

```
struct pam_conv {
    int (**conv)();
    void (**appdata_ptr);
}
```

The argument **conv** is defined as:

```
int conv( int num_msg, const struct pam_message **msg,
          const struct pam_response **resp, void *appdata );
```

The conversation function, **conv**, allows PAM to send messages to, and get input from, a user. The arguments to the function have the following definition and behavior:

### num\_msg

The number of lines of messages to be displayed (all messages are returned in one-line fragments, each no longer than **PAM\_MAX\_MSG\_SIZE** characters and with no more lines than **PAM\_MAX\_NUM\_MSG**)

**msg** Contains the message text and its style.

```
struct pam_message {
    int style; /* Message style */
    char *msg; /* The message */
}
```

The message style, can be one of:

### PAM\_PROMPT\_ECHO\_OFF

Prompts users with message and does not echo their responses; it is typically for use with requesting passwords and other sensitive information.

### PAM\_PROMPT\_ECHO\_ON

Prompts users with message and echoes their responses back to them.

### PAM\_ERROR\_MSG

Displays message as an error message.

### PAM\_TEXT\_INFO

Displays general information, such as authentication failures.

**resp** Holds the user's response and a response code.

```
struct pam_response {
    char **resp; /* Reference to the response */
    int resp_retcode; /* Not used, should be 0 */
}
```

### appdata, appdata\_ptr

Pointers to the application data that can be passed by the calling application to the PAM modules. Use these to allow PAM to send data back to the application.

## PAMHandle

The PAM handle representing the current user authentication session is returned upon successful completion.

## Return Values

Upon successful completion, **pam\_start** returns **PAM\_SUCCESS**, and a reference to the pointer of a valid PAM handle is returned through *PAMHandle*. If the routine fails, a value different from **PAM\_SUCCESS** is returned, and the *PAMHandle* reference is NULL.

## Error Codes

<b>PAM_SERVICE_ERR</b>	An error occurred in a PAM module.
<b>PAM_SYSTEM_ERR</b>	A system error occurred.
<b>PAM_BUF_ERR</b>	A memory error occurred.

## Related Information

“pam\_end Subroutine” on page 951, “pam\_set\_data Subroutine” on page 959, “pam\_set\_item Subroutine” on page 960

---

## pam\_strerror Subroutine

### Purpose

Translates a PAM error code to a string message.

### Library

PAM Library (**libpam.a**)

### Syntax

```
#include <security/pam_appl.h>

const char *pam_strerror (PAMHandle, ErrorCode)
pam_handle_t *PAMHandle;
int ErrorCode;
```

### Description

The **pam\_strerror** subroutine uses the error number returned by the PAM routines and returns the PAM error message that is associated with that error number. If the error number is not known to **pam\_strerror**, or there is no translation error message, then NULL is returned. The caller should not free or modify the returned string.

### Parameters

<i>PAMhandle</i>	The PAM handle representing the current user authentication session. This handle is obtained by a call to <b>pam_start()</b> .
<i>ErrorCode</i>	The PAM error code for which the PAM error message is to be retrieved.

### Return Values

Upon successful completion, **pam\_strerror** returns the PAM error message corresponding to the PAM error code, *ErrorCode*. A NULL pointer is returned if the routine fails, the error code is not known, or no error message exists for that error code.

---

## passwdexpired Subroutine

### Purpose

Checks the user’s password to determine if it has expired.

## Syntax

```
passwdexpired ( UserName, Message)  
char *UserName;  
char **Message;
```

## Description

The **passwdexpired** subroutine checks a user's password to determine if it has expired. The subroutine checks the **registry** variable in the **/etc/security/user** file to ascertain where the user is administered. If the **registry** variable is not defined, the **passwdexpired** subroutine checks the local, NIS, and DCE databases for the user definition and expiration time.

The **passwdexpired** subroutine may pass back informational messages, such as how many days remain until password expiration.

## Parameters

<i>UserName</i>	Specifies the user's name whose password is to be checked.
<i>Message</i>	Points to a pointer that the <b>passwdexpired</b> subroutine allocates memory for and fills in. This string is suitable for printing and issues messages, such as in how many days the password will expire.

## Return Values

Upon successful completion, the **passwdexpired** subroutine returns a value of 0. If this subroutine fails, it returns one of the following values:

- 1 Indicates that the password is expired, and the user must change it.
- 2 Indicates that the password is expired, and only a system administrator may change it.
- 1 Indicates that an internal error has occurred, such as a memory allocation (malloc) failure or database corruption.

## Error Codes

The **passwdexpired** subroutine fails if one or more of the following values is true:

<b>ENOENT</b>	Indicates that the user could not be found.
<b>EACCES</b>	Indicates that the user did not have permission to check password expiration.
<b>ENOMEM</b>	Indicates that memory allocation (malloc) failed.
<b>EINVAL</b>	Indicates that the parameters are not valid.

## Related Information

The **authenticate** ("authenticate Subroutine" on page 113) subroutine.

The **login** command.

---

## passwdexpiredx Subroutine

### Purpose

Checks the user's password to determine if it has expired, in multiple methods.

## Syntax

```
passwdexpiredx (UserName, Message, State)  
char *UserName;  
char **Message;  
char **State;
```

## Description

The **passwdexpiredx** subroutine checks a user's password to determine if it has expired. The subroutine uses the user's **SYSTEM** attribute to ascertain which administrative domains are used for password authentication.

The **passwdexpiredx** subroutine can pass back informational messages, such as how many days remain until password expiration.

The *State* parameter can contain information about the results of the authentication process. The *State* parameter from an earlier call to the **authenticatex** subroutine can be used to control how password expiration checking is performed. Authentication mechanisms that were not used to authenticate a user are not examined for expired passwords. The *State* parameter must be initialized to reference a null pointer if the *State* parameter from an earlier call to the **authenticatex** subroutine is not used.

## Parameters

<i>UserName</i>	Specifies the user's name whose password is to be checked.
<i>Message</i>	Points to a pointer that the <b>passwdexpiredx</b> subroutine allocates memory for and fills in. This string is suitable for printing, and it issues messages, such as an alert that indicates how many days are left before the password expires.
<i>State</i>	Points to a pointer that the <b>passwdexpiredx</b> subroutine allocates memory for and fills in. The <i>State</i> parameter can also be the result of an earlier call to the <b>authenticatex</b> subroutine. The <i>State</i> parameter contains information about the results of the password expiration examination process for each term in the user's <b>SYSTEM</b> attribute. The calling application is responsible for freeing this memory when it is no longer needed for a subsequent call to the <b>chpassx</b> subroutine.

## Return Values

Upon successful completion, the **passwdexpiredx** subroutine returns a value of 0. If this subroutine fails, it returns one of the following values:

-1	Indicates that an internal error has occurred, such as a memory allocation (malloc) failure or database corruption.
1	Indicates that one or more passwords are expired, and the user must change it. None of the expired passwords require system administrator intervention to be changed.
2	Indicates that one or more passwords are expired, at least one of which must be changed by the user and at least one of which requires system administrator intervention to be changed.
3	Indicates that all expired passwords require system administrator intervention to be changed.

## Error Codes

The **passwdexpiredx** subroutine fails if one or more of the following values is true:

<b>EACCES</b>	The user did not have permission to access the password attributes required to check password expiration.
<b>EINVAL</b>	The parameters are not valid.
<b>ENOENT</b>	The user could not be found.
<b>ENOMEM</b>	Memory allocation (malloc) failed.

## Related Information

The “`authenticate` Subroutine” on page 115.

The `login` Command.

---

## passwdpolicy Subroutine

### Purpose

Supports password strength policies on a per-user or per-named-policy basis.

### Syntax

```
#include <passwdpolicy.h>
int passwdpolicy (const char *name, int type, const char *old_password,
                 const char *new_password, time64_t last_update);
```

### Description

The `passwdpolicy` subroutine supports application use of password strength policies on a per-user or per-named-policy basis. The policies that are supported include password dictionaries, history list length, history list expiration, maximum lifetime of a password, minimum period of time between permitted password changes, maximum period after which an expired password can be changed, maximum number of repeated characters in a password, minimum number of alphabetic characters in a password, minimum number of nonalphabetic characters in a password, minimum length of a password, and a list of loadable modules that can be used to determine additional password strength rules.

The `type` parameter allows an application to select where the policy values are located. Privileged process can use either `PWP_USERNAME` or `PWP_SYSTEMPOLICY`. Unprivileged processes are limited to using `PWP_LOCALPOLICY`.

The following named attributes are used for each test:

<b>dictionlist</b>	A <b>SEC_LIST</b> value that gives a list of dictionaries to be checked. If <code>new_password</code> is found in any of the named dictionaries, this test fails. If this test fails, the return value contains the <b>PWP_IN_DICTIONARY</b> bit.
<b>histsize</b>	A <b>SEC_INT</b> value giving the permissible size of the named user's password history. The named user's password history is obtained by calling <code>getuserattr</code> with the <b>S_HISTLIST</b> attribute. If this attribute does not exist, password history checks are disabled. A value of 0 disables password history tests. If this test fails, the return value contains the <b>PWP_REUSED_PW</b> bit.
<b>histexpire</b>	A <b>SEC_INT</b> value that gives the number of weeks that must elapse before a password in the named user's password history list can be reused. If this test fails the return value contains the <b>PWP_REUSED_TOO_SOON</b> bit.
<b>maxage</b>	A <b>SEC_INT</b> value that gives the number of weeks a password can be considered valid. A password that has not been modified more recently than <b>maxage</b> weeks is considered to have expired and is subject to the <b>maxexpired</b> test. A value less than or equal to 0 disables this test. This attribute is used to determine if <b>maxexpired</b> must be tested, and it does not generate a return value.
<b>minage</b>	A <b>SEC_INT</b> value that gives the number of weeks before a password can be changed. A password that has been modified more recently than <b>minage</b> weeks fails this test. A value less than or equal to 0 disables this test. If this test fails, the return value contains the <b>PWP_TOO_SOON</b> bit.
<b>maxexpired</b>	A <b>SEC_INT</b> value that gives the number of weeks after which an expired password cannot be changed. A value of 0 indicates that an expired password cannot be changed. A value of -1 indicates that an expired password can be changed after any length of time. If this test fails, the return value contains the <b>PWP_EXPIRED</b> bit.



<b>maxrepeats</b>	A <b>SEC_INT</b> value that gives the maximum number of times any single character can appear in the new password. A value less than or equal to 0 disables this test. If this test fails, the return value contains the <b>PWP_TOO_MANY_REPEATS</b> bit.
<b>mindiff</b>	A <b>SEC_INT</b> value that gives the maximum number of characters in the new password that must not be present in the old password. A value less than or equal to 0 disables this test. If this test fails, the return value contains the <b>PWP_TOO_MANY_SAME</b> bit.
<b>minalpha</b>	A <b>SEC_INT</b> value that gives the minimum number of alphabetic characters that must be present in the password. A value less than or equal to 0 disables this test. If this test fails, the return value contains the <b>PWP_TOO_FEW_ALPHA</b> bit.
<b>minother</b>	A <b>SEC_INT</b> value that gives the minimum number of nonalphabetic characters that must be present in the password. A value less than or equal to 0 disables this test. If this test fails, the return value contains the bit <b>PWP_TOO_FEW_OTHER</b> bit.
<b>minlen</b>	A <b>SEC_INT</b> value that gives the minimum required length of a password. There is no maximum value for this attribute. A value less than or equal to 0 disables this test. If this test fails, the return value contains the <b>PWP_TOO_SHORT</b> bit.
<b>pwdchecks</b>	A <b>SEC_LIST</b> value that gives a list of named loadable modules that must be executed to validate the password. If this test fails, the return value contains the <b>PWP_FAILED_OTHER</b> bit.

## Parameters

<i>name</i>	The name of either a specific user or named policy. User names have policy information determined by invoking the <b>getuserattr</b> subroutine. Policy names have policy information determined by invoking the <b>getconfattr</b> subroutine.
<i>type</i>	One of three values:  <b>PWP_USERNAME</b> Policy values for <b>PWP_USERNAME</b> are stored in <b>/etc/security/user</b> . Password tests <b>PWP_REUSED_PW</b> and <b>PWP_REUSED_TOO_SOON</b> are only enabled for this value.  <b>PWP_SYSTEMPOLICY</b> Policy values for <b>PWP_SYSTEMPOLICY</b> are stored in <b>/etc/security/passwd_policy</b> .  <b>PWP_LOCALPOLICY</b> Policy values for <b>PWP_LOCALPOLICY</b> are stored in <b>/usr/lib/security/passwd_policy</b> .
<i>old_password</i>	The current value of the password. This function does not verify that <i>old_password</i> is the correct current password. Invoking <b>passwdpolicy</b> with a NULL pointer for this parameter disables <b>PWP_TOO_MANY_SAME</b> tests.
<i>new_password</i>	The value of the new password. Invoking <b>passwdpolicy</b> with a NULL pointer for this parameter disables all tests except password age tests.
<i>last_update</i>	The time the password was last changed, as a <b>time64_t</b> value, expressed in seconds since the UNIX® epoch. A 0 value for this parameter disables password age tests regardless of the value of any other parameter.

## Return Values

The return value is a bit-mapped representation of the tests that failed. A return value of 0 indicates that all password rules passed. A value of -1 indicates that some other error, other than a failed password test, has occurred. The **errno** variable indicates the cause of that error. Applications must compare a non-zero return value against -1 before checking any specific bits in the return value.

## Files

The **/usr/include/pwdpolicy.h** header file.

## Related Information

“passwdexpired Subroutine” on page 973, “passwdstrength Subroutine”

---

## passwdstrength Subroutine

### Purpose

Performs basic password age and construction tests.

### Syntax

```
#include <pwdpolicy.h>
int passwdstrength (const char *old_password, const char *new_password,
                  time64_t last_update, passwd_policy_t *policy, int checks);
```

### Description

The **passwdstrength** subroutine performs basic password age and construction tests. Password history, reuse, and dictionary tests are not performed. The values contained in the *policy* parameter are used to validate the value of *new\_password*.

The following fields are used by the **passwdstrength** subroutine.

<b>pwp_version</b>	Specifies the version of the <b>passwd_policy_t</b> structure. The current structure version number is <b>PWP_VERSION_1</b> .
<b>pwp_minage</b>	The number of seconds, as a <b>time32_t</b> , between the time a password is modified and the time the password can again be modified. This field is referenced if <b>PWP_TOO_SOON</b> is set in <i>checks</i> .
<b>pwp_maxage</b>	The number of seconds, as a <b>time32_t</b> , after which a password that has been modified is considered to be expired. This field is referenced if <b>PWP_EXPIRED</b> is set in <i>checks</i> .
<b>pwp_maxexpired</b>	The number of seconds, as a <b>time32_t</b> , since a password has expired after which it can no longer be modified. A value of 0 indicates that an expired password cannot be changed. A value of -1 indicates that an expired password can be changed after any length of time. This field is referenced if <b>PWP_EXPIRED</b> is set in <i>checks</i> .
<b>pwp_minalpha</b>	The minimum number of characters in the password that must be alphabetic characters, as determined by invoking the <b>isalpha()</b> macro. A value less than or equal to 0 disables this test. This field is referenced if <b>PWP_TOO_FEW_ALPHA</b> is set in <i>checks</i> .
<b>pwp_minother</b>	The minimum number of characters in the password that cannot be alphabetic characters, as determined by invoking the <b>isalpha()</b> macro. A value less than or equal to 0 disables this test. This field is referenced if <b>PWP_TOO_FEW_OTHER</b> is set in <i>checks</i> .
<b>pwp_minlen</b>	The minimum total number of characters in the password. A value less than or equal to 0 disables this test. This field is referenced if <b>PWP_TOO_SHORT</b> is set in <i>checks</i> .
<b>pwp_maxrepeats</b>	The maximum number of times an individual character can appear in the password. A value less than or equal to 0 disables this test. This field is referenced if <b>PWP_TOO_MANY_REPEATS</b> is set in <i>checks</i> .
<b>pwp_mindiff</b>	The minimum number of characters that must be changed between the old password and the new password. A value less than or equal to 0 disables this test. If this test fails, the return value contains the bit <b>PWP_TOO_MANY_SAME</b> . This field is referenced if <b>PWP_TOO_MANY_SAME</b> is set in <i>checks</i> .

### Parameters

<i>old_password</i>	The value of the current password. This parameter must be non-NULL if <b>PWP_TOO_MANY_SAME</b> is set in <i>checks</i> or the results are undefined.
---------------------	--

<i>new_password</i>	The value of the new password. This parameter must be non-NULL if any of <b>PWP_TOO_SHORT</b> , <b>PWP_TOO_FEW_ALPHA</b> , <b>PWP_TOO_FEW_OTHER</b> , <b>PWP_TOO_MANY_SAME</b> , or <b>PWP_TOO_MANY_REPEATS</b> are set in <i>checks</i> or the results are undefined.
<i>last_update</i>	The time the password was last changed, as a <b>time64_t</b> value, expressed in seconds since the UNIX epoch. A 0 value for this parameter indicates that the password has never been set. This might cause <b>PWP_EXPIRED</b> to be set in the return value if it is set in <i>checks</i> .
<i>policy</i>	A pointer to a <b>passwd_policy_t</b> containing the values for the password policy attributes.
<i>checks</i>	A bitmask value that indicates the set of password tests to be performed. The return value contains only those bits that are defined in <i>checks</i> .

## Return Values

The return value is a bit-mapped representation of the tests that failed. A return value of 0 indicates that all password rules requested in the *checks* parameter passed. A value of -1 indicates that some other error, other than a password test, has occurred. The **errno** variable indicates the cause of that error. Applications must compare a non-zero return value against -1 before checking any specific bits in the return value.

## Files

The `/usr/include/pwdpolicy.h` header file.

## Related Information

“passwdexpired Subroutine” on page 973, “passwdpolicy Subroutine” on page 976

---

## pathconf or fpathconf Subroutine

### Purpose

Retrieves file-implementation characteristics.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
long pathconf ( Path, Name)
const char *Path;
int Name;
```

```
long fpathconf( FileDescriptor, Name)
int FileDescriptor, Name;
```

### Description

The **pathconf** subroutine allows an application to determine the characteristics of operations supported by the file system contained by the file named by the *Path* parameter. Read, write, or execute permission of the named file is not required, but all directories in the path leading to the file must be searchable.

The **fpathconf** subroutine allows an application to retrieve the same information for an open file.

## Parameters

<i>Path</i>	Specifies the path name.
<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>Name</i>	Specifies the configuration attribute to be queried. If this attribute is not applicable to the file specified by the <i>Path</i> or <i>FileDescriptor</i> parameter, the <b>pathconf</b> subroutine returns an error. Symbolic values for the <i>Name</i> parameter are defined in the <b>unistd.h</b> file:
<b>_PC_LINK_MAX</b>	Specifies the maximum number of links to the file.
<b>_PC_MAX_CANON</b>	Specifies the maximum number of bytes in a canonical input line. This value is applicable only to terminal devices.
<b>_PC_MAX_INPUT</b>	Specifies the maximum number of bytes allowed in an input queue. This value is applicable only to terminal devices.
<b>_PC_NAME_MAX</b>	Specifies the maximum number of bytes in a file name, not including a terminating null character. This number can range from 14 through 255. This value is applicable only to a directory file.
<b>_PC_PATH_MAX</b>	Specifies the maximum number of bytes in a path name, including a terminating null character.
<b>_PC_PIPE_BUF</b>	Specifies the maximum number of bytes guaranteed to be written atomically. This value is applicable only to a first-in-first-out (FIFO).
<b>_PC_CHOWN_RESTRICTED</b>	Returns 0 if the use of the <b>chown</b> subroutine is restricted to a process with appropriate privileges, and if the <b>chown</b> subroutine is restricted to changing the group ID of a file only to the effective group ID of the process or to one of its supplementary group IDs.
<b>_PC_NO_TRUNC</b>	Returns 0 if long component names are truncated. This value is applicable only to a directory file.
<b>_PC_VDISABLE</b>	This is always 0. No disabling character is defined. This value is applicable only to a terminal device.
<b>_PC_AIX_DISK_PARTITION</b>	Determines the physical partition size of the disk. <b>Note:</b> The <b>_PC_AIX_DISK_PARTITION</b> variable is available only to the root user.
<b>_PC_AIX_DISK_SIZE</b>	Determines the disk size in megabytes. <b>Note:</b> The <b>_PC_AIX_DISK_SIZE</b> variable is available only to the root user.
<b>Note:</b> The <b>_PC_FILESIZEBITS</b> and <b>PC_SYNC_IO</b> flags apply to AIX 4.3 and later releases.	
<b>_PC_FILESIZEBITS</b>	Returns the minimum number of bits required to hold the file system's maximum file size as a signed integer. The smallest value returned is <b>32</b> .
<b>_PC_SYNC_IO</b>	Returns <b>-1</b> if the file system does not support the <b>Synchronized Input and Output</b> option. Any value other than <b>-1</b> is returned if the file system supports the option.

## Notes:

1. If the *Name* parameter has a value of **\_PC\_LINK\_MAX**, and if the *Path* or *FileDescriptor* parameter refers to a directory, the value returned applies to the directory itself.
2. If the *Name* parameter has a value of **\_PC\_NAME\_MAX** or **\_PC\_NO\_TRUNC**, and if the *Path* or *FileDescriptor* parameter refers to a directory, the value returned applies to filenames within the directory.
3. If the *Name* parameter has a value of **\_PC\_PATH\_MAX**, and if the *Path* or *FileDescriptor* parameter refers to a directory that is the working directory, the value returned is the maximum length of a relative pathname.
4. If the *Name* parameter has a value of **\_PC\_PIPE\_BUF**, and if the *Path* parameter refers to a FIFO special file or the *FileDescriptor* parameter refers to a pipe or a FIFO special file, the value returned applies to the referenced object. If the *Path* or *FileDescriptor* parameter refers to a directory, the value returned applies to any FIFO special file that exists or can be created within the directory.
5. If the *Name* parameter has a value of **\_PC\_CHOWN\_RESTRICTED**, and if the *Path* or *FileDescriptor* parameter refers to a directory, the value returned applies to any files, other than directories, that exist or can be created within the directory.

## Return Values

If the **pathconf** or **fpathconf** subroutine is successful, the specified parameter is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error. If the variable corresponding to the *Name* parameter has no limit for the *Path* parameter or the *FileDescriptor* parameter, both the **pathconf** and **fpathconf** subroutines return a value of -1 without changing the **errno** global variable.

## Error Codes

The **pathconf** or **fpathconf** subroutine fails if the following error occurs:

**EINVAL**                The name parameter specifies an unknown or inapplicable characteristic.

The **pathconf** subroutine can also fail if any of the following errors occur:

<b>EACCES</b>	Search permission is denied for a component of the path prefix.
<b>EINVAL</b>	The implementation does not support an association of the <i>Name</i> parameter with the specified file.
<b>ENAMETOOLONG</b>	The length of the <i>Path</i> parameter string exceeds the <b>PATH_MAX</b> value.
<b>ENAMETOOLONG</b>	<i>Pathname</i> resolution of a symbolic link produced an intermediate result whose length exceeds <b>PATH_MAX</b> .
<b>ENOENT</b>	The named file does not exist or the <i>Path</i> parameter points to an empty string.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>ELOOP</b>	Too many symbolic links were encountered in resolving path.

The **fpathconf** subroutine can fail if either of the following errors occur:

<b>EBADF</b>	The <i>File Descriptor</i> parameter is not valid.
<b>EINVAL</b>	The implementation does not support an association of the <i>Name</i> parameter with the specified file.

## Related Information

The “chown, fchown, lchown, chownx, or fchownx Subroutine” on page 151, “confstr Subroutine” on page 181, **sysconf** subroutine.

Files, Directories, and File Systems for Programmers, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pause Subroutine

### Purpose

Suspends a process until a signal is received.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
int pause (void)
```

### Description

The **pause** subroutine suspends the calling process until it receives a signal. The signal must not be one that is ignored by the calling process. The **pause** subroutine does not affect the action taken upon the receipt of a signal.

### Return Values

If the signal received causes the calling process to end, the **pause** subroutine does not return.

If the signal is caught by the calling process and control is returned from the signal-catching function, the calling process resumes execution from the point of suspension. The **pause** subroutine returns a value of -1 and sets the **errno** global variable to **EINTR**.

### Related Information

The **incinterval**, **alarm**, or **settimer** (“getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386) subroutine, **kill** or **killpg** (“kill or killpg Subroutine” on page 580) subroutine, **sigaction**, **sigvec**, or **signal** subroutine, **wait**, **waitpid**, or **wait3** subroutine.

---

## pcap\_close Subroutine

### Purpose

Closes the open files related to the packet capture descriptor and frees the memory used by the packet capture descriptor.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>

void pcap_close(pcap_t * p);
```

### Description

The **pcap\_close** subroutine closes the files associated with the packet capture descriptor and deallocates resources. If the **pcap\_open\_offline** subroutine was previously called, the **pcap\_close** subroutine closes the *savefile*, a previously saved packet capture data file. Or the **pcap\_close** subroutine closes the packet capture device if the **pcap\_open\_live** subroutine was previously called.

## Parameters

*p* Points to a packet capture descriptor as returned by the **pcap\_open\_live** or the **pcap\_open\_offline** subroutine.

## Related Information

The **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_compile Subroutine

### Purpose

Compiles a filter expression into a filter program.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_compile(pcap_t * p, struct bpf_program *fp, char * str,  
int optimize, bpf_u_int32 netmask);
```

### Description

The **pcap\_compile** subroutine is used to compile the string *str* into a filter program. This filter program will then be used to filter, or select, the desired packets.

### Parameters

<i>netmask</i>	Specifies the <i>netmask</i> of the network device. The <i>netmask</i> can be obtained from the <b>pcap_lookupnet</b> subroutine.
<i>optimize</i>	Controls whether optimization on the resulting code is performed.
<i>p</i>	Points to a packet capture descriptor returned from the <b>pcap_open_offline</b> or the <b>pcap_open_live</b> subroutine.
<i>program</i>	Points to a <b>bpf_program</b> struct which will be filled in by the <b>pcap_compile</b> subroutine if the subroutine is successful.
<i>str</i>	Contains the filter expression.

### Return Values

Upon successful completion, the **pcap\_compile** subroutine returns 0, and the program parameter will hold the filter program. If **pcap\_compile** subroutine is unsuccessful, -1 is returned.

### Related Information

The **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_lookupnet** (“pcap\_lookupnet Subroutine” on page 991) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine, **pcap\_setfilter** (“pcap\_setfilter Subroutine” on page 999) subroutine.

---

## pcap\_datalink Subroutine

### Purpose

Obtains the link layer type for the packet capture device.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_datalink(pcap_t * p);
```

### Description

The **pcap\_datalink** subroutine returns the link layer type of the packet capture device, for example, IFT\_ETHER. This is useful in determining the size of the datalink header at the beginning of each packet that is read.

### Parameters

*p* Points to the packet capture descriptor as returned by the **pcap\_open\_live** or the **pcap\_open\_offline** subroutine.

### Return Values

The **pcap\_datalink** subroutine returns the link layer type.

**Note:** Only call this subroutine after successful calls to either the **pcap\_open\_live** or the **pcap\_open\_offline** subroutine. Never call the **pcap\_datalink** subroutine after a call to **pcap\_close** as unpredictable results will occur.

### Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_dispatch Subroutine

### Purpose

Collects and processes packets.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_dispatch(pcap_t * p, int cnt, pcap_handler callback,  
u_char * user);
```



## Description

The **pcap\_dispatch** subroutine reads and processes packets. This subroutine can be called to read and process packets that are stored in a previously saved packet capture data file, known as the *savefile*. The subroutine can also read and process packets that are being captured live.

Notice that the third parameter, *callback*, is of the type **pcap\_handler**. This is a pointer to a user-provided subroutine with three parameters. Define this user-provided subroutine as follows:

```
void user_routine(u_char *user, struct pcap_pkthdr *phdr, u_char *pdata)
```

The parameter, *user*, is the *user* parameter that is passed into the **pcap\_dispatch** subroutine. The parameter, *phdr*, is a pointer to the **pcap\_pkthdr** structure which precedes each packet in the *savefile*. The parameter, *pdata*, points to the packet data. This allows users to define their own handling of packet capture data.

## Parameters

<i>callback</i>	Points to a user-provided routine that will be called for each packet read. The user is responsible for providing a valid pointer, and that unpredictable results can occur if an invalid pointer is supplied. <b>Note:</b> The <b>pcap_dump</b> subroutine can also be specified as the <i>callback</i> parameter. If this is done, the <b>pcap_dump_open</b> subroutine should be called first. The pointer to the <b>pcap_dumper_t</b> struct returned from the <b>pcap_dump_open</b> subroutine should be used as the <i>user</i> parameter to the <b>pcap_dispatch</b> subroutine. The following program fragment illustrates this use: <pre>pcap_dumper_t *pd pcap_t * p; int rc = 0;  pd = pcap_dump_open(p, "/tmp/savefile");  rc = pcap_dispatch(p, 0 , pcap_dump, (u_char *) pd);</pre>
<i>cnt</i>	Specifies the maximum number of packets to process before returning. A <i>cnt</i> of -1 processes all the packets received in one buffer. A <i>cnt</i> of 0 processes all packets until an error occurs, EOF is reached, or the read times out (when doing live reads and a non-zero read timeout is specified).
<i>p</i>	Points to a packet capture descriptor returned from the <b>pcap_open_offline</b> or the <b>pcap_open_live</b> subroutine. This will be used to store packet data that is read in.
<i>user</i>	Specifies the first argument to pass into the <i>callback</i> routine.

## Return Values

Upon successful completion, the **pcap\_dispatch** subroutine returns the number of packets read. If EOF is reached in a *savefile*, zero is returned. If the **pcap\_dispatch** subroutine is unsuccessful, -1 is returned. In this case, the **pcap\_geterr** or **pcap\_perror** subroutine can be used to get the error text.

## Related Information

The **pcap\_dump** (“pcap\_dump Subroutine” on page 986) subroutine, **pcap\_dump\_close** (“pcap\_dump\_close Subroutine” on page 986) subroutine, **pcap\_dump\_open** (“pcap\_dump\_open Subroutine” on page 987) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine.

---

## pcap\_dump Subroutine

### Purpose

Writes packet capture data to a binary file.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
void pcap_dump(u_char * user, struct pcap_pkthdr * h, u_char * sp);
```

### Description

The **pcap\_dump** subroutine writes the packet capture data to a binary file. The packet header data, contained in *h*, will be written to the file pointed to by the *user* file pointer, followed by the packet data from *sp*. Up to *h->caplen* bytes of *sp* will be written.

The file that *user* points to (where the **pcap\_dump** subroutine writes to) must be open. To open the file and retrieve its pointer, use the **pcap\_dump\_open** subroutine.

The calling arguments for the **pcap\_dump** subroutine are suitable for use with **pcap\_dispatch** subroutine and the **pcap\_loop** subroutine. To retrieve this data, the **pcap\_open\_offline** subroutine can be invoked with the name of the file that *user* points to as its first parameter.

### Parameters

<i>h</i>	Contains the packet header data that will be written to the packet capture data file, known as the <i>savefile</i> . This data will be written ahead of the rest of the packet data.
<i>sp</i>	Points to the packet data that is to be written to the <i>savefile</i> .
<i>user</i>	Specifies the <i>savefile</i> file pointer which is returned from the <b>pcap_dump_open</b> subroutine. It should be cast to a <code>u_char *</code> when passed in.

### Related Information

The **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984) subroutine, **pcap\_dump\_close** (“pcap\_dump\_close Subroutine”) subroutine, **pcap\_dump\_open** (“pcap\_dump\_open Subroutine” on page 987) subroutine, **pcap\_loop** (“pcap\_loop Subroutine” on page 992) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_dump\_close Subroutine

### Purpose

Closes a packet capture data file, known as a *savefile*.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
void pcap_dump_close(pcap_dumper_t * p);
```

## Description

The **pcap\_dump\_close** subroutine closes a packet capture data file, known as the *savefile*, that was opened using the **pcap\_dump\_open** subroutine.

## Parameters

*p* Points to a **pcap\_dumper\_t**, which is synonymous with a FILE \*, the file pointer of a *savefile*.

## Related Information

The **pcap\_dump\_open** (“pcap\_dump\_open Subroutine”) subroutine.

---

## pcap\_dump\_open Subroutine

### Purpose

Opens or creates a file for writing packet capture data.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
pcap_dumper_t *pcap_dump_open(pcap_t * p, char * fname);
```

### Description

The **pcap\_dump\_open** subroutine opens or creates the packet capture data file, known as the *savefile*. This action is specified through the *fname* parameter. The subroutine then writes the required packet capture file header to the file. The **pcap\_dump** subroutine can then be called to write the packet capture data associated with the packet capture descriptor, **p**, into this file. The **pcap\_dump\_open** subroutine must be called before calling the **pcap\_dump** subroutine.

### Parameters

*fname* Specifies the name of the file to open. A “-” indicates that standard output should be used instead of a file.

*p* Specifies a packet capture descriptor returned by the **pcap\_open\_offline** or the **pcap\_open\_live** subroutine.

### Return Values

Upon successful completion, the **pcap\_dump\_open** subroutine returns a pointer to the file that was opened or created. This pointer is a pointer to a **pcap\_dumper\_t**, which is synonymous with FILE \*. See the **pcap\_dump** (“pcap\_dump Subroutine” on page 986), **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984), or the **pcap\_loop** (“pcap\_loop Subroutine” on page 992) subroutine for an example of how to

use **pcap\_dumper\_t**. If the **pcap\_dump\_open** subroutine is unsuccessful, Null is returned. Use the **pcap\_geterr** subroutine to obtain the specific error text.

## Related Information

The **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984) subroutine, **pcap\_dump** (“pcap\_dump Subroutine” on page 986) subroutine, **pcap\_dump\_close** (“pcap\_dump\_close Subroutine” on page 986) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_loop** (“pcap\_loop Subroutine” on page 992) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_file Subroutine

### Purpose

Obtains the file pointer to the *savefile*, a previously saved packed capture data file.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
FILE *pcap_file(pcap_t * p);
```

### Description

The **pcap\_file** subroutine returns the file pointer to the *savefile*. If there is no open *savefile*, 0 is returned. This subroutine should be called after a successful call to the **pcap\_open\_offline** subroutine and before any calls to the **pcap\_close** subroutine.

### Parameters

<i>p</i>	Points to a packet capture descriptor as returned by the <b>pcap_open_offline</b> subroutine.
----------	---

### Return Values

The **pcap\_file** subroutine returns the file pointer to the *savefile*.

### Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_fileno Subroutine

### Purpose

Obtains the descriptor for the packet capture device.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
int pcap_fileno(pcap_t * p);
```

## Description

The **pcap\_fileno** subroutine returns the descriptor for the packet capture device. This subroutine should be called only after a successful call to the **pcap\_open\_live** subroutine and before any calls to the **pcap\_close** subroutine.

## Parameters

*p* Points to a packet capture descriptor as returned by the **pcap\_open\_live** subroutine.

## Return Values

The **pcap\_fileno** subroutine returns the descriptor for the packet capture device.

## Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine.

---

## pcap\_geterr Subroutine

### Purpose

Obtains the most recent pcap error message.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
char *pcap_geterr(pcap_t * p);
```

## Description

The **pcap\_geterr** subroutine returns the error text pertaining to the last pcap library error. This subroutine is useful in obtaining error text from those subroutines that do not return an error string. Since the pointer returned points to a memory space that will be reused by the pcap library subroutines, it is important to copy this message into a new buffer if the error text needs to be saved.

## Parameters

*p* Points to a packet capture descriptor as returned by the **pcap\_open\_live** or the **pcap\_open\_offline** subroutine.

## Return Values

The **pcap\_geterr** subroutine returns a pointer to the most recent error message from a pcap library subroutine. If there were no previous error messages, a string with 0 as the first byte is returned.

## Related Information

The **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine, **pcap\_strerror** (“pcap\_strerror Subroutine” on page 1001) subroutine.

---

## pcap\_is\_swapped Subroutine

### Purpose

Reports if the byte order of the previously saved packet capture data file, known as the *savefile* was swapped.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_is_swapped(pcap_t * p);
```

### Description

The **pcap\_is\_swapped** subroutine returns 1 (True) if the current *savefile* uses a different byte order than the current system. This subroutine should be called after a successful call to the **pcap\_open\_offline** subroutine and before any calls to the **pcap\_close** subroutine.

### Parameters

<i>p</i>	Points to a packet capture descriptor as returned from the <b>pcap_open_offline</b> subroutine.
----------	---

### Return Values

1	If the byte order of the <i>savefile</i> is different from that of the current system.
0	If the byte order of the <i>savefile</i> is the same as that of the current system.

## Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_lookupdev Subroutine

### Purpose

Obtains the name of a network device on the system.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
char *pcap_lookupdev(char * errbuf);
```

## Description

The **pcap\_lookupdev** subroutine gets a network device suitable for use with the **pcap\_open\_live** and the **pcap\_lookupnet** subroutines. If no interface can be found, or none are configured to be up, Null is returned. In the case of multiple network devices attached to the system, the **pcap\_lookupdev** subroutine returns the first one it finds to be up, other than the loopback interface. (Loopback is always ignored.)

## Parameters

*errbuf*

Returns error text and is only set when the **pcap\_lookupdev** subroutine fails.

## Return Values

Upon successful completion, the **pcap\_lookupdev** subroutine returns a pointer to the name of a network device attached to the system. If **pcap\_lookupdev** subroutine is unsuccessful, Null is returned, and text indicating the specific error is written to *errbuf*.

## Related Information

The **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_lookupnet** (“pcap\_lookupnet Subroutine”) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine.

---

## pcap\_lookupnet Subroutine

### Purpose

Returns the network address and subnet mask for a network device.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
int pcap_lookupnet(char * device, bpf_u_int32 * netp, bpf_u_int32 * maskp,  
char * errbuf);
```

## Description

Use the **pcap\_lookupnet** subroutine to determine the network address and subnet mask for the network device, **device**.

## Parameters

*device*

Specifies the name of the network device to use for the network lookup, for example, en0.

*errbuf*

Returns error text and is only set when the **pcap\_lookupnet** subroutine fails.

*maskp*  
*netp*

Holds the subnet mask associated with **device**.  
Holds the network address for the **device**.

## Return Values

Upon successful completion, the **pcap\_lookupnet** subroutine returns 0. If the **pcap\_lookupnet** subroutine is unsuccessful, -1 is returned, and *errbuf* is filled in with an appropriate error message.

## Related Information

The **pcap\_compile** (“pcap\_compile Subroutine” on page 983) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_lookupdev** (“pcap\_lookupdev Subroutine” on page 990) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine.

---

## pcap\_loop Subroutine

### Purpose

Collects and processes packets.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_loop(pcap_t * p, int cnt, pcap_handler callback,  
              u_char * user);
```

### Description

The **pcap\_loop** subroutine reads and processes packets. This subroutine can be called to read and process packets that are stored in a previously saved packet capture data file, known as the *savefile*. The subroutine can also read and process packets that are being captured live.

This subroutine is similar to **pcap\_dispatch** subroutine except it continues to read packets until *cnt* packets have been processed, EOF is reached (in the case of offline reading), or an error occurs. It does not return when live read timeouts occur. That is, specifying a non-zero read timeout to the **pcap\_open\_live** subroutine and then calling the **pcap\_loop** subroutine allows the reception and processing of any packets that arrive when the timeout occurs.

Notice that the third parameter, *callback*, is of the type **pcap\_handler**. This is a pointer to a user-provided subroutine with three parameters. Define this user-provided subroutine as follows:

```
void user_routine(u_char *user, struct pcap_pkthdr *phdr, u_char *pdata)
```

The parameter, *user*, will be the user parameter that was passed into the **pcap\_dispatch** subroutine. The parameter, *phdr*, is a pointer to the **pcap\_pkthdr** structure, which precedes each packet in the *savefile*. The parameter, *pdata*, points to the packet data. This allows users to define their own handling of their filtered packets.



## Parameters

<i>callback</i>	<p>Points to a user-provided routine that will be called for each packet read. The user is responsible for providing a valid pointer, and that unpredictable results can occur if an invalid pointer is supplied.</p> <p><b>Note:</b> The <b>pcap_dump</b> subroutine can also be specified as the callback parameter. If this is done, call the <b>pcap_dump_open</b> subroutine first. Then use the pointer to the <b>pcap_dumper_t</b> struct returned from the <b>pcap_dump_open</b> subroutine as the user parameter to the <b>pcap_dispatch</b> subroutine. The following program fragment illustrates this use:</p> <pre>pcap_dumper_t *pd pcap_t * p; int rc = 0;  pd = pcap_dump_open(p, "/tmp/savefile");  rc = pcap_dispatch(p, 0 , pcap_dump, (u_char *) pd);</pre>
<i>cnt</i>	<p>Specifies the maximum number of packets to process before returning. A negative value causes the <b>pcap_loop</b> subroutine to loop forever, or until EOF is reached or an error occurs. A <i>cnt</i> of 0 processes all packets until an error occurs or EOF is reached.</p>
<i>p</i>	<p>Points to a packet capture descriptor returned from the <b>pcap_open_offline</b> or the <b>pcap_open_live</b> subroutine. This will be used to store packet data that is read in.</p>
<i>user</i>	<p>Specifies the first argument to pass into the <i>callback</i> routine.</p>

## Return Values

Upon successful completion, the **pcap\_loop** subroutine returns 0. 0 is also returned if EOF has been reached in a *savefile*. If the **pcap\_loop** subroutine is unsuccessful, -1 is returned. In this case, the **pcap\_geterr** subroutine or the **pcap\_perror** subroutine can be used to get the error text.

## Related Information

The **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984) subroutine, **pcap\_dump** (“pcap\_dump Subroutine” on page 986) subroutine, **pcap\_dump\_close** (“pcap\_dump\_close Subroutine” on page 986) subroutine, **pcap\_dump\_open** (“pcap\_dump\_open Subroutine” on page 987) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine.

---

## pcap\_major\_version Subroutine

### Purpose

Obtains the major version number of the packet capture format used to write the *savefile*, a previously saved packet capture data file.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_major_version(pcap_t * p);
```

## Description

The **pcap\_major\_version** subroutine returns the major version number of the packet capture format used to write the *savefile*. If there is no open *savefile*, 0 is returned.

**Note:** This subroutine should be called only after a successful call to **pcap\_open\_offline** subroutine and before any calls to the **pcap\_close** subroutine.

## Parameters

*p* Points to a packet capture descriptor as returned from **pcap\_open\_offline** subroutine.

## Return Values

The major version number of the packet capture format used to write the *savefile*.

## Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_minor\_version Subroutine

### Purpose

Obtains the minor version number of the packet capture format used to write the *savefile*.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_minor_version(pcap_t * p);
```

### Description

The **pcap\_minor\_version** subroutine returns the minor version number of the packet capture format used to write the *savefile*. This subroutine should only be called after a successful call to the **pcap\_open\_offline** subroutine and before any calls to the **pcap\_close** subroutine.

### Parameters

*p* Points to a packet capture descriptor as returned from the **pcap\_open\_offline** subroutine.

### Return Values

The minor version number of the packet capture format used to write the *savefile*.

### Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_next Subroutine

### Purpose

Obtains the next packet from the packet capture device.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
u_char *pcap_next(pcap_t * p, struct pcap_pkthdr * h);
```

### Description

The **pcap\_next** subroutine returns a `u_char` pointer to the next packet from the packet capture device. The packet capture device can be a network device or a *savefile* that contains packet capture data. The data has the same format as used by **tcpdump**.

### Parameters

*h*

Points to the packet header of the packet that is returned. This is filled in upon return by this routine.

*p*

Points to the packet capture descriptor to use as returned by the **pcap\_open\_live** or the **pcap\_open\_offline** subroutine.

### Return Values

Upon successful completion, the **pcap\_next** subroutine returns a pointer to a buffer containing the next packet and fills in the *h*, which points to the packet header of the returned packet. If the **pcap\_next** subroutine is unsuccessful, Null is returned.

### Related Information

The **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984) subroutine, **pcap\_dump** (“pcap\_dump Subroutine” on page 986) subroutine, **pcap\_dump\_close** (“pcap\_dump\_close Subroutine” on page 986) subroutine, **pcap\_dump\_open** (“pcap\_dump\_open Subroutine” on page 987) subroutine, **pcap\_loop** (“pcap\_loop Subroutine” on page 992) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine”) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

The **tcpdump** command.

---

## pcap\_open\_live Subroutine

### Purpose

Opens a network device for packet capture.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
pcap_t *pcap_open_live( const char * device, const int snaplen,  
const int promisc, const int to_ms, char * ebuf);
```

## Description

The **pcap\_open\_live** subroutine opens the specified network device for packet capture. The term "live" is to indicate that a network device is being opened, as opposed to a file that contains packet capture data. This subroutine must be called before any packet capturing can occur. All other routines dealing with packet capture require the packet capture descriptor that is created and initialized with this routine. See the **pcap\_open\_offline** ("pcap\_open\_offline Subroutine" on page 997) subroutine for more details on opening a previously saved file that contains packet capture data.

## Parameters

<i>device</i>	Specifies a string that contains the name of the network device to open for packet capture, for example, en0.
<i>ebuf</i>	Returns error text and is only set when the <b>pcap_open_live</b> subroutine fails.
<i>promisc</i>	Specifies that the device is to be put into promiscuous mode. A value of 1 (True) turns promiscuous mode on. If this parameter is 0 (False), the device will remain unchanged. In this case, if it has already been set to promiscuous mode (for some other reason), it will remain in this mode.
<i>snaplen</i>	Specifies the maximum number of bytes to capture per packet.
<i>to_ms</i>	Specifies the read timeout in milliseconds.

## Return Values

Upon successful completion, the **pcap\_open\_live** subroutine will return a pointer to the packet capture descriptor that was created. If the **pcap\_open\_live** subroutine is unsuccessful, Null is returned, and text indicating the specific error is written into the *ebuf* buffer.

## Related Information

The **pcap\_close** ("pcap\_close Subroutine" on page 982) subroutine, **pcap\_compile** ("pcap\_compile Subroutine" on page 983) subroutine, **pcap\_datalink** ("pcap\_datalink Subroutine" on page 984) subroutine, **pcap\_dispatch** ("pcap\_dispatch Subroutine" on page 984) subroutine, **pcap\_dump** ("pcap\_dump Subroutine" on page 986) subroutine, **pcap\_dump\_open** ("pcap\_dump\_open Subroutine" on page 987) subroutine, **pcap\_geterr** ("pcap\_geterr Subroutine" on page 989) subroutine, **pcap\_loop** ("pcap\_loop Subroutine" on page 992) subroutine, **pcap\_open\_offline** ("pcap\_open\_offline Subroutine" on page 997) subroutine, **pcap\_perror** ("pcap\_perror Subroutine" on page 998) subroutine, **pcap\_setfilter** ("pcap\_setfilter Subroutine" on page 999) subroutine, **pcap\_snapshot** ("pcap\_setfilter Subroutine" on page 999) subroutine, **pcap\_stats** ("pcap\_stats Subroutine" on page 1000) subroutine.

---

## pcap\_open\_live\_sb Subroutine

### Purpose

Opens a network device for packet capture, allowing you to specify the buffer length of a Berkeley Packet Filter (BPF).

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
pcap_t * pcap_open_live_sb( const char *device, int snaplen,
int promisc, int to_ms, char *ebuf, int buflen )
```

## Description

The **pcap\_open\_live\_sb** subroutine opens the specified network device for packet capture. This subroutine allows you to specify the buffer size for the BPF to use in capturing the packets. You must run this subroutine before any packet capturing can occur. All other subroutines dealing with packet capture require the packet capture descriptor that is created and initialized with this subroutine.

To opening a previously saved file that contains packet capture data, use the **pcap\_open\_offline** subroutine.

## Parameters

<i>buf_len</i>	Specifies the buffer size that the BPF is to use. If the system cannot provide memory of this size, the system will choose a smaller size.
<i>device</i>	Specifies a string that contains the name of the network device to open for packet capture, for example, en0.
<i>ebuf</i>	Returns error text and is only set when the <b>pcap_open_live</b> subroutine fails.
<i>promisc</i>	Specifies that the device is to be put into the promiscuous mode. A value of 1 (True) turns the promiscuous mode on. If this parameter is zero (False), the device remains unchanged. In this case, if it has already been set to the promiscuous mode (for some other reason), it remains in this mode.
<i>snaplen</i>	Specifies the maximum number of bytes to capture per packet.
<i>to_ms</i>	Specifies the read timeout in milliseconds.

## Return Values

If successful, the **pcap\_open\_live\_sb** subroutine returns a pointer to the packet capture descriptor that is created. If the **pcap\_open\_live\_sb** subroutine is unsuccessful, NULL is returned, and the text indicating the specific error is written into the *ebuf* buffer.

## Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984) subroutine, **pcap\_file** (“pcap\_file Subroutine” on page 988) subroutine, **pcap\_fileno** (“pcap\_fileno Subroutine” on page 988) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_is\_swapped** (“pcap\_is\_swapped Subroutine” on page 990) subroutine, **pcap\_loop** (“pcap\_loop Subroutine” on page 992) subroutine, **pcap\_major\_version** (“pcap\_major\_version Subroutine” on page 993) subroutine, **pcap\_minor\_version** (“pcap\_minor\_version Subroutine” on page 994) subroutine, **pcap\_next** (“pcap\_next Subroutine” on page 995) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, and **pcap\_open\_offline** (“pcap\_open\_offline Subroutine”) subroutine.

---

## pcap\_open\_offline Subroutine

### Purpose

Opens a previously saved file containing packet capture data.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
pcap_t *pcap_open_offline(char * fname, char * ebuf);
```

## Description

The **pcap\_open\_offline** subroutine opens a previously saved packet capture data file, known as the *savefile*. This subroutine creates and initializes a packet capture (pcap) descriptor and opens the specified *savefile* containing the packet capture data for reading.

This subroutine should be called before any other related routines that require a packet capture descriptor for offline packet processing. See the **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine for more details on live packet capture.

**Note:** The format of the *savefile* is expected to be the same as the format used by the **tcpdump** command.

## Parameters

*ebuf*

Returns error text and is only set when the **pcap\_open\_offline** subroutine fails.

*fname*

Specifies the name of the file to open. A hyphen (-) passed as the *fname* parameter indicates that stdin should be used as the file to open.

## Return Values

Upon successful completion, the **pcap\_open\_offline** subroutine will return a pointer to the newly created packet capture descriptor. If the **pcap\_open\_offline** subroutine is unsuccessful, Null is returned, and text indicating the specific error is written into the *ebuf* buffer.

## Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_dispatch** (“pcap\_dispatch Subroutine” on page 984) subroutine, **pcap\_file** (“pcap\_file Subroutine” on page 988) subroutine, **pcap\_fileno** (“pcap\_fileno Subroutine” on page 988) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_is\_swapped** (“pcap\_is\_swapped Subroutine” on page 990) subroutine, **pcap\_loop** (“pcap\_loop Subroutine” on page 992) subroutine, **pcap\_major\_version** (“pcap\_major\_version Subroutine” on page 993) subroutine, **pcap\_minor\_version** (“pcap\_minor\_version Subroutine” on page 994) subroutine, **pcap\_next** (“pcap\_next Subroutine” on page 995) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine.

The **tcpdump** command.

---

## pcap\_perror Subroutine

### Purpose

Prints the passed-in prefix, followed by the most recent error text.

### Library

pcap Library (**libpcap.a**)

## Syntax

```
#include <pcap.h>
```

```
void pcap_perror(pcap_t * p, char * prefix);
```

## Description

The **pcap\_perror** subroutine prints the text of the last pcap library error to stderr, prefixed by *prefix*. If there were no previous errors, only *prefix* is printed.

## Parameters

<i>p</i>	Points to a packet capture descriptor as returned by the <b>pcap_open_live</b> subroutine or the <b>pcap_open_offline</b> subroutine.
<i>prefix</i>	Specifies the string that is to be printed before the stored error message.

## Related Information

The **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine, **pcap\_strerror** (“pcap\_strerror Subroutine” on page 1001) subroutine.

---

## pcap\_setfilter Subroutine

### Purpose

Loads a filter program into a packet capture device.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_setfilter(pcap_t * p, struct bpf_program * fp);
```

### Description

The **pcap\_setfilter** subroutine is used to load a filter program into the packet capture device. This causes the capture of the packets defined by the filter to begin.

### Parameters

<i>fp</i>	Points to a filter program as returned from the <b>pcap_compile</b> subroutine.
<i>p</i>	Points to a packet capture descriptor returned from the <b>pcap_open_offline</b> or the <b>pcap_open_live</b> subroutine.

### Return Values

Upon successful completion, the **pcap\_setfilter** subroutine returns 0. If the **pcap\_setfilter** subroutine is unsuccessful, -1 is returned. In this case, the **pcap\_geterr** subroutine can be used to get the error text, and the **pcap\_perror** subroutine can be used to display the text.

## Related Information

The **pcap\_compile** (“pcap\_compile Subroutine” on page 983) subroutine, **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine.

---

## pcap\_snapshot Subroutine

### Purpose

Obtains the number of bytes that will be saved for each packet captured.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>
```

```
int pcap_snapshot( pcap_t * p );
```

### Description

The **pcap\_snapshot** subroutine returns the snapshot length, which is the number of bytes to save for each packet captured.

**Note:** This subroutine should only be called after successful calls to either the **pcap\_open\_live** subroutine or **pcap\_open\_offline** subroutine. It should not be called after a call to the **pcap\_close** subroutine.

### Parameters

*p* Points to the packet capture descriptor as returned by the **pcap\_open\_live** or the **pcap\_open\_offline** subroutine.

### Return Values

The **pcap\_snapshot** subroutine returns the snapshot length.

## Related Information

The **pcap\_close** (“pcap\_close Subroutine” on page 982) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_open\_offline** (“pcap\_open\_offline Subroutine” on page 997) subroutine.

---

## pcap\_stats Subroutine

### Purpose

Obtains packet capture statistics.

### Library

pcap Library (**libpcap.a**)



## Syntax

```
#include <pcap.h>

int pcap_stats (pcap_t *p, struct pcap_stat *ps);
```

## Description

The **pcap\_stats** subroutine fills in a **pcap\_stat** struct. The values represent packet statistics from the start of the run to the time of the call. Statistics for both packets that are received by the filter and packets that are dropped are stored inside a **pcap\_stat** struct. This subroutine is for use when a packet capture device is opened using the **pcap\_open\_live** subroutine.

## Parameters

<i>p</i>	Points to a packet capture descriptor as returned by the <b>pcap_open_live</b> subroutine.
<i>ps</i>	Points to the <b>pcap_stat</b> struct that will be filled in with the packet capture statistics.

## Return Values

On successful completion, the **pcap\_stats** subroutine fills in *ps* and returns 0. If the **pcap\_stats** subroutine is unsuccessful, -1 is returned. In this case, the error text can be obtained with the **pcap\_perror** subroutine or the **pcap\_geterr** subroutine.

## Related Information

The **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_open\_live** (“pcap\_open\_live Subroutine” on page 995) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine.

---

## pcap\_strerror Subroutine

### Purpose

Obtains the error message indexed by *error*.

### Library

pcap Library (**libpcap.a**)

### Syntax

```
#include <pcap.h>

char *pcap_strerror(int error);
```

### Description

Lookup the error message indexed by *error*. The possible values of *error* correspond to the values of the *errno* global variable. This function is equivalent to the **strerror** subroutine.

### Parameters

<i>error</i>	Specifies the key to use in obtaining the corresponding error message. The error message is taken from the system's <b>sys_errlist</b> .
--------------	--

## Return Values

The **pcap\_strerror** subroutine returns the appropriate error message from the system error list.

## Related Information

The **pcap\_geterr** (“pcap\_geterr Subroutine” on page 989) subroutine, **pcap\_perror** (“pcap\_perror Subroutine” on page 998) subroutine, **strerror** subroutine.

---

## pclose Subroutine

### Purpose

Closes a pipe to a process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
int pclose ( Stream)
FILE *Stream;
```

### Description

The **pclose** subroutine closes a pipe between the calling program and a shell command to be executed. Use the **pclose** subroutine to close any stream you opened with the **popen** subroutine. The **pclose** subroutine waits for the associated process to end, and then returns the exit status of the command.

**Attention:** If the original processes and the **popen** process are reading or writing a common file, neither the **popen** subroutine nor the **pclose** subroutine should use buffered I/O. If they do, the results are unpredictable.

Avoid problems with an output filter by flushing the buffer with the **fflush** subroutine.

### Parameter

*Stream* Specifies the **FILE** pointer of an opened pipe.

### Return Values

The **pclose** subroutine returns a value of -1 if the *Stream* parameter is not associated with a **popen** command or if the status of the child process could not be obtained. Otherwise, the value of the termination status of the command language interpreter is returned; this will be 127 if the command language interpreter cannot be executed.

### Error Codes

If the application has:

- Called the **wait** subroutine,
- Called the **waitpid** subroutine with a process ID less than or equal to zero or equal to the process ID of the command line interpreter,
- Masked the SIGCHLD signal, or
- Called any other function that could perform one of the steps above, and

one of these calls caused the termination status to be unavailable to the **pclose** subroutine, a value of -1 is returned and the **errno** global variable is set to **ECHILD**.

## Related Information

The **fclose** or **fflush** (“fclose or fflush Subroutine” on page 253) subroutine, **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **pipe** (“pipe Subroutine” on page 1025) subroutine, **popen** (“popen Subroutine” on page 1144) subroutine, **wait**, **waitpid**, or **wait3** subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## perfstat\_cpu Subroutine

### Purpose

Retrieves individual logical processor usage statistics.

### Library

perfstat library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_cpu (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t * name;
perfstat_cpu_t * userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

The **perfstat\_cpu** subroutine retrieves one or more individual processor usage statistics. The same function can be used to retrieve the number of available sets of logical processor statistics.

To get one or more sets of processor usage metrics, set the *name* parameter to the name of the first processor for which statistics are desired, and set the *desired\_number* parameter. To start from the first processor, set the *name* parameter to *""*. The *userbuff* parameter must always point to a memory area big enough to contain the desired number of **perfstat\_cpu\_t** structures that will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next processor, or to *""* after all structures have been copied.

To retrieve the number of available sets of processor usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The returned value will be the number of available sets.

This number represents the number of logical processors for which statistics are available. In a dynamic LPAR environment, this number is the highest logical index of an online processor since the last reboot. See the Perfstat API article in Performance Tools and APIs Technical Reference for more information on the **perfstat\_cpu** subroutine and DLPAR.

In AIX 5.3 and later, SPLPAR environments virtualize physical processors. To help accurately measure the resource use in a virtualized environment, the POWER5 family of processors implements a register PURR (Processor Utilization Resource Register) for each core. The PURR is a 64-bit counter with the same units as the timebase register and tracks the real physical processor resource used on a per-thread or

per-partition level. The PURR registers are not compatible with previous global counters (user, system, idle and wait fields) returned by the **perfstat\_cpu** and the **perfstat\_cpu\_total** subroutines. All data consumers requiring processor utilization must be modified to support PURR-based computations as shown in the example for `perfstat_partition_total()` interface under Perfstat API programming.

## Parameters

<i>name</i>	Contains either "", FIRST_CPU, or a name identifying the first logical processor for which statistics are desired. Logical processor names are:  cpu0, cpu1,...
<i>userbuff</i>	To provide binary compatibility with previous versions of the library, names like <i>proc0</i> , <i>proc1</i> , ... will still be accepted. These names will be treated as if their corresponding <i>cpuN</i> name was used, but the names returned in the structures will always be names starting with <i>cpu</i> .
<i>sizeof_struct</i>	Points to the memory area that is to be filled with one or more <b>perfstat_cpu_t</b> structures.
<i>desired_number</i>	Specifies the size of the <b>perfstat_cpu_t</b> structure: <code>sizeof(perfstat_cpu_t)</code> .
	Specifies the number of <b>perfstat_cpu_t</b> structures to copy to <i>userbuff</i> .

## Return Values

Unless the **perfstat\_cpu** subroutine is used to retrieve the number of available structures, the number of structures filled is returned upon successful completion. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_cpu** subroutine is unsuccessful if the following is true:

**EINVAL**                      One of the parameters is not valid.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

The “perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu\_total Subroutine,” “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_cpu\_total Subroutine

### Purpose

Retrieves global processor usage statistics.

### Library

Perfstat library (**libperfstat.a**)

## Syntax

```
#include <libperfstat.h>
```

```
int perfstat_cpu_total (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_cpu_total_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

## Description

The **perfstat\_cpu\_total** subroutine returns global processor usage statistics in a **perfstat\_cpu\_total\_t** structure.

To get statistics that are global to the whole system, the *name* parameter must be set to NULL, the *userbuff* parameter must be allocated, and the *desired\_number* parameter must be set to 1.

The **perfstat\_cpu\_total** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

In AIX 5.3 and later, SPLPAR environments virtualize physical processors. To help accurately measure the resource used in a virtualized environment, the POWER5 family of processors implements a register PURR (Processor Utilization Resource Register) for each core. The PURR is a 64-bit counter with the same units as the timebase register and tracks the real physical processor resource used on a per-thread or per-partition level. The PURR registers are not compatible with previous global counters (user, system, idle and wait fields) returned by the **perfstat\_cpu** and the **perfstat\_cpu\_total** subroutines. All data consumers requiring processor use must be modified to support PURR-based computations as shown in the example for the **perfstat\_partition\_total** interface under Perfstat API programming.

## Parameters

<i>name</i>	Must set to NULL.
<i>userbuff</i>	Points to the memory area that is to be filled with the <b>perfstat_cpu_total_t</b> structure.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_cpu_total_t</b> structure: <code>sizeof(perfstat_cpu_total_t)</code> .
<i>desired_number</i>	Must set to 1.

## Return Values

Upon successful completion, the number of structures filled is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_cpu\_total** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.
<b>ENOMEM</b>	The string default length is too short.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_disk Subroutine,” “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingospace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, and “perfstat\_protocol Subroutine” on page 1022.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_disk Subroutine

### Purpose

Retrieves individual disk usage statistics.

### Library

Perfstat library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_disk (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_disk_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

The **perfstat\_disk** subroutine retrieves one or more individual disk usage statistics. The same function can also be used to retrieve the number of available sets of disk statistics.

To get one or more sets of disk usage metrics, set the *name* parameter to the name of the first disk for which statistics are desired, and set the *desired\_number* parameter. To start from the first disk, specify "" or FIRST\_DISK as the *name*. The *userbuff* parameter must always point to a memory area big enough to contain the desired number of **perfstat\_disk\_t** structures that will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next disk, or to "" after all structures have been copied.

To retrieve the number of available sets of disk usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The returned value will be the number of available sets.

The **perfstat\_disk** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

To improve system performance, the collection of disk input and output statistics is disabled by default in current releases of AIX.

To enable the collection of this data, run:

```
chdev -l sys0 -a iostat=true
```

To display the current setting, run:

```
lsattr -E -l sys0 -a iostat
```

Another way to enable the collection of the disk input and output statistics is to use the **sys\_parm** API and the **SYSP\_V\_IOSTRUN** flag:

To get the current status of the flag, run the following:

```
struct vario var;
sys_parm(SYSP_GET,SYSP_V_IOSTRUN, &var);
```

To set the flag, run the following:

```
struct vario var;
var.v.v_iostrun.value=1; /* 1 to set & 0 to unset */
sys_parm(SYSP_SET,SYSP_V_IOSTRUN, &var);
```

## Parameters

<i>name</i>	Contains either "", FIRST_DISK, or a name identifying the first disk for which statistics are desired. For example:  hdisk0, hdisk1, ...
<i>userbuff</i>	Points to the memory area to be filled with one or more <b>perfstat_disk_t</b> structures.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_disk_t</b> structure: sizeof(perfstat_disk_t)
<i>desired_number</i>	Specifies the number of <b>perfstat_disk_t</b> structures to copy to <i>userbuff</i> .

## Return Values

Unless the function is used to retrieve the number of available structures, the number of structures filled is returned upon successful completion. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_disk** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.
<b>ENOMEM</b>	The string default length is too short.
<b>ENOMSG</b>	Cannot access the dictionary.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_diskadapter Subroutine

### Purpose

Retrieves individual disk adapter usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_diskadapter (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_diskadapter_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

The **perfstat\_diskadapter** subroutine retrieves one or more individual disk adapter usage statistics. The same function can be used to retrieve the number of available sets of adapter statistics.

To get one or more sets of disk adapter usage metrics, set the *name* parameter to the name of the first disk adapter for which statistics are desired, and set the *desired\_number* parameter. To start from the first disk adapter, set the *name* parameter to "" or FIRST\_DISKADAPTER. The *userbuff* parameter must point to a memory area big enough to contain the desired number of **perfstat\_diskadapter\_t** structures which will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next disk adapter, or to "" if all structures have been copied.

To retrieve the number of available sets of disk adapter usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The returned value will be the number of available sets.

The **perfstat\_diskadapter** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

To improve system performance, the collection of disk input/output statistics is disabled by default in current releases of AIX.

To enable the collection of this data, use:

```
chdev -l sys0 -a iostat=true
```

To display the current setting, use:

```
lsattr -E -l sys0 -a iostat
```

Another way to enable the collection of the disk input/output statistics is to use the **sys\_parm** API and the **SYSP\_V\_IOSTRUN** flag:

To get the current status of the flag:

```
struct vario var;
sys_parm(SYSP_GET,SYSP_V_IOSTRUN, &var);
```

To set the flag:



```

struct vario var;
var.v.v_iostrun.value=1; /* 1 to set & 0 to unset */
sys_parm(SYSP_SET,SYSP_V_IOSTRUN, &var);

```

## Parameters

<i>name</i>	Contains either "", FIRST_DISKADAPTER, or a name identifying the first disk adapter for which statistics are desired. For example:  scsi0, scsi1, ...
<i>userbuff</i>	Points to the memory area to be filled with one or more <b>perfstat_diskadapter_t</b> structures.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_diskadapter_t</b> structure: sizeof(perfstat_diskadapter_t)
<i>desired_number</i>	Specifies the number of <b>perfstat_diskadapter_t</b> structures to copy to <i>userbuff</i> .

## Return Values

Unless the function is used to retrieve the number of available structures, the number of structures filled is returned upon successful completion. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_diskadapter** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.
<b>ENOMEM</b>	The string default length is too short.
<b>ENOMSG</b>	Cannot access the dictionary.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskpath Subroutine,” “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_diskpath Subroutine

### Purpose

Retrieves individual disk path usage statistics.

### Library

Perfstat library (**libperfstat.a**)

## Syntax

```
#include <libperfstat.h>
```

```
int perfstat_diskpath (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_diskpath_t *userbuff
size_t sizeof_struct;
int desired_number;
```

## Description

The **perfstat\_diskpath** subroutine retrieves one or more individual disk path usage statistics. The same function can also be used to retrieve the number of available sets of disk path statistics.

To get one or more sets of disk path usage metrics, set the *name* parameter to the name of the first disk path for which statistics are desired, and set the *desired\_number* parameter. To start from the first disk path, specify "" or FIRST\_DISKPATH as the *name* parameter. To start from the first path of a specific disk, set the *name* parameter to the diskname. The *userbuff* parameter must always point to a memory area big enough to contain the desired number of **perfstat\_diskpath\_t** structures that will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next disk path, or to "" after all structures have been copied.

To retrieve the number of available sets of disk path usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The number of available sets is returned.

The **perfstat\_diskpath** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

To improve system performance, the collection of disk input and output statistics is disabled by default in current releases of AIX.

To enable the collection of this data, run:

```
chdev -l sys0 -a iostat=true
```

To display the current setting, run:

```
lsattr -E -l sys0 -a iostat
```

Another way to enable the collection of the disk input and output statistics is to use the **sys\_parm** API and the **SYSP\_V\_IOSTRUN** flag:

To get the current status of the flag, run the following:

```
struct vario var;
sys_parm(SYSP_GET,SYSP_V_IOSTRUN, &var);
```

To set the flag, run the following:

```
struct vario var;
var.v.v_iostrun.value=1; /* 1 to set & 0 to unset */
sys_parm(SYSP_SET,SYSP_V_IOSTRUN, &var);
```

## Parameters

<i>name</i>	Contains either "", FIRST_DISKPATH, a name identifying the first disk path for which statistics are desired, or a name identifying a disk for which path statistics are desired. For example: hdisk0_Path2, hdisk1_Path0, ... or hdisk5 (equivalent to hdisk5_Pathfirstpath)
<i>userbuff</i>	Points to the memory area to be filled with one or more <b>perfstat_diskpath_t</b> structures.

*sizeof\_struct* Specifies the size of the **perfstat\_diskpath\_t** structure: `sizeof(perfstat_diskpath_t)`  
*desired\_number* Specifies the number of **perfstat\_diskpath\_t** structures to copy to *userbuff*.

## Return Values

Unless the function is used to retrieve the number of available structures, the number of structures filled is returned upon successful completion. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_diskpath** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.
<b>ENOMEM</b>	The string default length is too short.
<b>ENOMSG</b>	Cannot access the dictionary.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine,” “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_disk\_total Subroutine

### Purpose

Retrieves global disk usage statistics.

### Library

Perfstat library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_disk_total (name, userbuff, sizeof_struct, desired_number)  
perfstat_id_t *name;  
perfstat_disk_total_t *userbuff;  
size_t sizeof_struct;  
int desired_number;
```

### Description

The **perfstat\_disk\_total** subroutine returns global disk usage statistics in a **perfstat\_disk\_total\_t** structure.

To get statistics that are global to the whole system, the *name* parameter must be set to NULL, the *userbuff* parameter must be allocated, and the *desired\_number* parameter must be set to 1.

The **perfstat\_disk\_total** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

To improve system performance, the collection of disk input and output statistics is disabled by default in current releases of AIX.

To enable the collection of this data, run:

```
chdev -l sys0 -a iostat=true
```

To display the current setting, run:

```
lsattr -E -l sys0 -a iostat
```

Another way to enable the collection of the disk input and output statistics is to use the **sys\_parm** API and the **SYSP\_V\_IOSTRUN** flag:

To get the current status of the flag, run the following:

```
struct vario var;
sys_parm(SYSP_GET,SYSP_V_IOSTRUN, &var);
```

To set the flag, run the following:

```
struct vario var;
var.v.v_iostrun.value=1; /* 1 to set & 0 to unset */
sys_parm(SYSP_SET,SYSP_V_IOSTRUN, &var);
```

## Parameters

<i>name</i>	Must set to NULL.
<i>userbuff</i>	Points to the memory area that is to be filled with one or more <b>perfstat_disk_total_t</b> structures.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_disk_total_t</b> structure: <code>sizeof(perfstat_disk_total_t)</code>
<i>desired_number</i>	Must set to 1.

## Return Values

Upon successful completion, the number of structures that could be filled is returned. This is always 1. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_disk\_total** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.
<b>ENOMEM</b>	The string default length is too short.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_memory\_total Subroutine,” “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingospace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_memory\_total Subroutine

### Purpose

Retrieves global memory usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_memory_total (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_memory_total_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

The **perfstat\_memory\_total** subroutine returns global memory usage statistics in a **perfstat\_memory\_total\_t** structure.

To get statistics that are global to the whole system, the *name* parameter must be set to NULL, the *userbuff* parameter must be allocated, and the *desired\_number* parameter must be set to 1.

### Parameters

<i>name</i>	Must be set to NULL.
<i>userbuff</i>	Points to the memory area that is to be filled with the <b>perfstat_memory_total_t</b> structure.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_memory_total_t</b> structure: sizeof(perfstat_memory_total_t).
<i>desired_number</i>	Must be set to 1.

### Return Values

Upon successful completion, the number of structures filled is returned. This will always be 1. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

### Error Codes

The **perfstat\_memory\_total** subroutine is unsuccessful if the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
---------------	-------------------------------------

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine,” “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, and “perfstat\_protocol Subroutine” on page 1022.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_netbuffer Subroutine

### Purpose

Retrieves network buffer allocation usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_netbuffer (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_netbuffer_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

The **perfstat\_netbuffer** subroutine retrieves statistics about network buffer allocations for each possible buffer size. Returned counts are the sum of allocation statistics for all processors (kernel statistics are kept per size per processor) corresponding to a buffer size.

To get one or more sets of network buffer allocation usage metrics, set the *name* parameter to the network buffer size for which statistics are desired, and set the *desired\_number* parameter. To start from the first network buffer size, specify "" or FIRST\_NETBUFFER in the *name* parameter. The *userbuff* parameter must point to a memory area big enough to contain the desired number of **perfstat\_netbuffer\_t** structures which will be copied by this function.

Upon return, the *name* parameter will be set to either the ASCII size of the next buffer type, or to "" if all structures have been copied. Only the statistics for network buffer sizes that have been used are returned. Consequently, there can be holes in the returned array of statistics, and the structure corresponding to allocations of size 4096 may directly follow the structure for size 256 (in case 512, 1024 and 2048 have not been used yet). The structure corresponding to a buffer size not used yet is returned (with all fields set to 0) when it is directly asked for by name.

To retrieve the number of available sets of network buffer usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The returned value will be the number of available sets.

## Parameters

<i>name</i>	Contains either "", FIRST_NETBUFFER, or the size of the network buffer in ASCII. It is a power of 2. For example: 32, 64, 128, ..., 16384
<i>userbuff</i>	Points to the memory area to be filled with one or more <b>perfstat_netbuffer_t</b> structures.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_netbuffer_t</b> structure: sizeof(perfstat_netbuffer_t)
<i>desired_number</i>	Specifies the number of <b>perfstat_netbuffer_t</b> structures to copy to <i>userbuff</i> .

## Return Values

Upon successful completion, the number of structures which could be filled is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_netbuffer** subroutine is unsuccessful if the following is true:

**EINVAL** One of the parameters is not valid.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_pagingospace Subroutine” on page 1018.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_netinterface Subroutine

### Purpose

Retrieves individual network interface usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_netinterface (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_netinterface_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

## Description

The **perfstat\_netinterface** subroutine retrieves one or more individual network interface usage statistics. The same function can also be used to retrieve the number of available sets of network interface statistics.

To get one or more sets of network interface usage metrics, set the *name* parameter to the name of the first network interface for which statistics are desired, and set the *desired\_number* parameter. To start from the first network interface, set the *name* parameter to "" or FIRST\_NETINTERFACE. The *userbuff* parameter must always point to a memory area big enough to contain the desired number of **perfstat\_netinterface\_t** structures that will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next network interface, or to "" after all structures have been copied.

To retrieve the number of available sets of network interface usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The returned value will be the number of available sets.

The **perfstat\_netinterface** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

## Parameters

<i>name</i>	Contains either "", FIRST_NETINTERFACE, or a name identifying the first network interface for which statistics are desired. For example:  en0, tr10, ...
<i>userbuff</i>	Points to the memory area that is to be filled with one or more <b>perfstat_netinterface_t</b> structures.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_netinterface_t</b> structure: sizeof(perfstat_netinterface_t)
<i>desired_number</i>	Specifies the number of <b>perfstat_netinterface_t</b> structures to copy to <i>userbuff</i> .

## Return Values

Upon successful completion unless the function is used to retrieve the number of available structures, the number of structures filled is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_netinterface** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.
<b>ENOMEM</b>	The string default length is too short.
<b>ENOMSG</b>	Cannot access the dictionary.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.



## perfstat\_netinterface\_total Subroutine

### Purpose

Retrieves global network interface usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_netinterface_total (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_netinterface_total_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

The **perfstat\_netinterface\_total** subroutine returns global network interface usage statistics in a **perfstat\_netinterface\_total\_t** structure.

To get statistics that are global to the whole system, the *name* parameter must be set to NULL, the *userbuff* parameter must be allocated, and the *desired\_number* parameter must be set to 1.

The **perfstat\_netinterface\_total** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

### Parameters

<i>name</i>	Must be set to NULL.
<i>userbuff</i>	Points to the memory area that is to be filled with the <b>perfstat_netinterface_total_t</b> structure.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_netinterface_total_t</b> structure: sizeof(perfstat_netinterface_total_t).
<i>desired_number</i>	Must be set to 1.

### Return Values

Upon successful completion, the number of structures filled is returned. This will always be 1. If unsuccessful, a value of -1 is returned and the **errno** variable is set.

### Error Codes

The **perfstat\_netinterface\_total** subroutine is unsuccessful if one of the following is true:

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.

### Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_pagingspace Subroutine,” “perfstat\_partial\_reset Subroutine” on page 1019, “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_pagingspace Subroutine

### Purpose

Retrieves individual paging space usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_pagingspace (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_pagingspace_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

### Description

This function retrieves one or more individual pagingspace usage statistics. The same functions can also be used to retrieve the number of available sets of paging space statistics.

To get one or more sets of paging space usage metrics, set the *name* parameter to the name of the first paging space for which statistics are desired, and set the *desired\_number* parameter. To start from the first paging space, set the *name* parameter to "" or FIRST\_PAGINGSPACE. In either case, *userbuff* must point to a memory area big enough to contain the desired number of **perfstat\_pagingspace\_t** structures which will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next paging space, or to "" if all structures have been copied.

To retrieve the number of available sets of paging space usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The number of available sets will be returned.

The **perfstat\_pagingspace** subroutine retrieves information from the ODM database. This information is automatically cached into a dictionary which is assumed to be frozen once loaded. The **perfstat\_reset** subroutine must be called to flush the dictionary whenever the machine configuration has changed.

### Parameters

<i>name</i>	Contains either "", FIRST_PAGINGSPACE, or a name identifying the first paging space for which statistics are desired. For example: paging00, hd6, ...
<i>userbuff</i>	Points to the memory area to be filled with one or more <b>perfstat_pagingspace_t</b> structures.

*sizeof\_struct* Specifies the size of the **perfstat\_pagingspace\_t** structure:  
`sizeof(perfstat_pagingspace_t)`

*desired\_number* Specifies the number of **perfstat\_pagingspace\_t** structures to copy to *userbuff*.

## Return Values

Unless the **perfstat\_pagingspace** subroutine is used to retrieve the number of available structures, the number of structures which could be filled is returned upon successful completion. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_pagingspace** subroutine is unsuccessful if one of the following are true:

**EINVAL** One of the parameters is not valid.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_partial\_reset Subroutine,” “perfstat\_protocol Subroutine” on page 1022, and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_partial\_reset Subroutine

### Purpose

Empties part of the libperfstat configuration information cache or resets system minimum and maximum counters for disks.

### Library

perfstat library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_partial_reset (name, resetmask)  
char * name;  
u_longlong_t resetmask;
```

### Description

The **perfstat\_cpu\_total**, **perfstat\_disk**, **perfstat\_diskadapter**, **perfstat\_netinterface**, and **perfstat\_pagingspace** subroutines return configuration information that is retrieved from the ODM database and automatically cached by the library. Other metrics provided by the LVM library and the **swapqry** subroutine are also cached for performance purpose.

The `perfstat_partial_reset` subroutine flushes some of this information cache and should be called whenever an identified part of the machine configuration has changed.

The `perfstat_partial_reset` subroutine can be used to reset a particular component (such as `hdisk0` or `en1`) when the `name` parameter is not NULL and the `resetmask` parameter contains only one bit. It can also be used to remove a whole category (such as disks or disk paths) from the cached information.

When the `name` parameter is NULL, the `resetmask` parameter can contain a combination of bits, such as `FLUSH_DISK|RESET_DISK_MINMAX|FLUSH_CPUTOTAL`. For more information on the `perfstat_partial_reset` subroutine, see `Perfstat API Programming`.

Several bit masks are available for the `resetmask` parameter. The behavior of the function is as follows:

<i>resetmask</i> value	Action when <i>name</i> is NULL	Action when <i>name</i> is not NULL and a single <i>resetmask</i> is set
FLUSH_CPUTOTAL	Flush <i>speed</i> and <i>description</i> in the <code>perfstat_cputotal_t</code> structure	An error is returned, and <code>errno</code> is set to <code>EINVAL</code> .
FLUSH_DISK	Flush <i>description</i> , <i>adapter</i> , <i>size</i> , <i>free</i> , and <i>vgname</i> in every <code>perfstat_disk_t</code> structure. Flush the list of disk adapters. Flush <i>size</i> , <i>free</i> , and <i>description</i> in every <code>perfstat_diskadapter_t</code> structure.	Flush <i>description</i> , <i>adapter</i> , <i>size</i> , <i>free</i> , and <i>vgname</i> in the specified <code>perfstat_disk_t</code> structure. Flush <i>adapter</i> in every <code>perfstat_diskpath_t</code> that matches the disk name followed by <code>_Path</code> . Flush <i>size</i> , <i>free</i> , and <i>description</i> of each <code>perfstat_diskadapter_t</code> that is linked to a path leading to this disk or to the disk itself.
RESET_DISK_ALL	Reset system resident all fields in every <code>perfstat_disk_t</code> structure.	An error is returned, and <code>errno</code> is set to <code>EINVAL</code> .
RESET_DISK_MINMAX	Reset system resident <i>min_rserv</i> , <i>max_rserv</i> , <i>min_wserv</i> , <i>max_wserv</i> , <i>wq_min_time</i> and <i>wq_max_time</i> in every <code>perfstat_disk_t</code> structure.	An error is returned, and <code>errno</code> is set to <code>ENOTSUP</code> .
FLUSH_DISKADAPTER	Flush the list of disk adapters. Flush <i>size</i> , <i>free</i> , and <i>description</i> in every <code>perfstat_diskadapter_t</code> structure. Flush <i>adapter</i> in every <code>perfstat_diskpath_t</code> structure. Flush <i>description</i> and <i>adapter</i> in every <code>perfstat_disk_t</code> structure.	Flush the list of disk adapters. Flush <i>size</i> , <i>free</i> , and <i>description</i> in the specified <code>perfstat_diskadapter_t</code> structure.
FLUSH_DISKPATH	Flush <i>adapter</i> in every <code>perfstat_diskpath_t</code> structure.	Flush <i>adapter</i> in the specified <code>perfstat_diskpath_t</code> structure.
FLUSH_PAGINGSPEACE	Flush the list of paging spaces. Flush <i>automatic</i> , <i>type</i> , <i>lpsize</i> , <i>mbsize</i> , <i>hostname</i> , <i>filename</i> , and <i>vgname</i> in every <code>perfstat_pagingspace_t</code> structure.	Flush the list of paging spaces. Flush <i>automatic</i> , <i>type</i> , <i>lpsize</i> , <i>mbsize</i> , <i>hostname</i> , <i>filename</i> , and <i>vgname</i> in the specified <code>perfstat_pagingspace_t</code> structure.
FLUSH_NETINTERFACE	Flush <i>description</i> in every <code>perfstat_netinterface_t</code> structure.	Flush <i>description</i> in the specified <code>perfstat_netinterface_t</code> structure.

## Parameters

- name* Contains a name identifying the component that metrics should be reset from the libperfstat cache. If this parameter is NULL, matches every component.
- resetmask* The category of the component if the *name* parameter is not NULL. The available values are listed in the preceding table. In case the *name* parameter is NULL, the *resetmask* parameter can be a combination of bits.

## Return Values

The `perfstat_partial_reset` subroutine returns a value of 0 upon successful completion. If unsuccessful, a value of -1 is returned, and the `errno` global variable is set to the appropriate code.

## Error Codes

**EINVAL** One of the parameters is not valid.

## Files

The `libperfstat.h` file defines standard macros, data types, and subroutines.

## Related Information

The “`perfstat_cpu` Subroutine” on page 1003, “`perfstat_cpu_total` Subroutine” on page 1004, “`perfstat_disk` Subroutine” on page 1006, “`perfstat_diskadapter` Subroutine” on page 1008, “`perfstat_diskpath` Subroutine” on page 1009, “`perfstat_disk_total` Subroutine” on page 1011, “`perfstat_memory_total` Subroutine” on page 1013, “`perfstat_netbuffer` Subroutine” on page 1014, “`perfstat_netinterface` Subroutine” on page 1015, “`perfstat_netinterface_total` Subroutine” on page 1017, “`perfstat_pagingspace` Subroutine” on page 1018, “`perfstat_partition_total` Subroutine,” “`perfstat_protocol` Subroutine” on page 1022, and “`perfstat_reset` Subroutine” on page 1023.

Perfstat API Programming.

---

## perfstat\_partition\_total Subroutine

### Purpose

Retrieves global Micro-Partitioning usage statistics.

### Library

perfstat library (`libperfstat.a`)

### Syntax

```
#include <libperfstat.h>
int perfstat_partition_total(name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_partition_total_t *userbuff;
size_t sizeof_struct;
int desired_number;
u_longlong_t reserved_pages;
u_longlong_t reserved_pagesize.
```

### Description

The `perfstat_partition_total` subroutine returns global Micro-Partitioning usage statistics in a `perfstat_partition_total_t` structure. To retrieve statistics that are global to the whole system, the `name` parameter must be set to NULL, the `userbuff` parameter must be allocated, and the `desired_number` parameter must be set to 1.

### Parameters

<i>name</i>	Must be set to NULL.
<i>userbuff</i>	Points to the memory area to be filled with the <code>perfstat_partition_total_t</code> structures.

<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_partition_total_t</b> structure: <code>sizeof(perfstat_partition_total_t)</code> .
<i>desired_number</i>	Must be set to 1.
<i>reserved_pagesize</i>	Specifies the size of the pages for reserved memory. Not for use with <b>DR</b> operations.
<i>reserved_pages</i>	Specifies the number of pages of type <i>reserved_pagesize</i> . This information can be retrieved by calling <b>vmgetinfo</b> . Not for use with <b>DR</b> operations.

## Return Values

Upon successful completion, the number of structures filled is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

<b>EINVAL</b>	One of the parameters is not valid.
<b>EFAULT</b>	Insufficient memory.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

“perfstat\_cpu Subroutine” on page 1003, “perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_disk\_total Subroutine” on page 1011, “perfstat\_memory\_total Subroutine” on page 1013, “perfstat\_netbuffer Subroutine” on page 1014, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_netinterface\_total Subroutine” on page 1017, “perfstat\_pagingspace Subroutine” on page 1018, “perfstat\_protocol Subroutine,” and “perfstat\_reset Subroutine” on page 1023.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_protocol Subroutine

### Purpose

Retrieves protocol usage statistics.

### Library

Perfstat Library (**libperfstat.a**)

### Syntax

```
#include <libperfstat.h>
```

```
int perfstat_protocol (name, userbuff, sizeof_struct, desired_number)
perfstat_id_t *name;
perfstat_protocol_t *userbuff;
size_t sizeof_struct;
int desired_number;
```

## Description

The **perfstat\_protocol** subroutine retrieves protocol usage statistics such as ICMP, ICMPv6, IP, IPv6, TCP, UDP, RPC, NFS, NFSv2, NFSv3. To get one or more sets of protocol usage metrics, set the *name* parameter to the name of the first protocol for which statistics are desired, and set the *desired\_number* parameter.

To start from the first protocol, set the *name* parameter to "" or FIRST\_PROTOCOL. The *userbuff* parameter must point to a memory area big enough to contain the desired number of **perfstat\_protocol\_t** structures which will be copied by this function. Upon return, the *name* parameter will be set to either the name of the next protocol, or to "" if all structures have been copied.

To retrieve the number of available sets of protocol usage metrics, set the *name* and *userbuff* parameters to NULL, and the *desired\_number* parameter to 0. The returned value will be the number of available sets.

## Parameters

<i>name</i>	Contains either "ip", "ipv6", "icmp", "icmpv6", "tcp", "udp", "rpc", "nfs", "nfsv2", "nfsv3", "", or FIRST_PROTOCOL.
<i>userbuff</i>	Points to the memory area to be filled with one or more <b>perfstat_protocol_t</b> structures.
<i>sizeof_struct</i>	Specifies the size of the <b>perfstat_protocol_t</b> structure: sizeof(perfstat_protocol_t)
<i>desired_number</i>	Specifies the number of <b>perfstat_protocol_t</b> structures to copy to <i>userbuff</i> .

## Return Values

Upon successful completion, the number of structures which could be filled is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set.

## Error Codes

The **perfstat\_protocol** subroutine is unsuccessful if the following is true:

**EINVAL**            One of the parameters is not valid.

## Files

The **libperfstat.h** file defines standard macros, data types, and subroutines.

## Related Information

"perfstat\_netbuffer Subroutine" on page 1014, "perfstat\_cpu Subroutine" on page 1003, "perfstat\_cpu\_total Subroutine" on page 1004, "perfstat\_disk Subroutine" on page 1006, "perfstat\_diskadapter Subroutine" on page 1008, "perfstat\_diskpath Subroutine" on page 1009, "perfstat\_disk\_total Subroutine" on page 1011, "perfstat\_memory\_total Subroutine" on page 1013, "perfstat\_netinterface Subroutine" on page 1015, "perfstat\_netinterface\_total Subroutine" on page 1017, "perfstat\_pagingspace Subroutine" on page 1018, and "perfstat\_partial\_reset Subroutine" on page 1019.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perfstat\_reset Subroutine

### Purpose

Empties libperfstat configuration information cache.

## Library

Perfstat Library (**libperfstat.a**)

## Syntax

```
#include <libperfstat.h>

void perfstat_reset (void)
```

## Description

The **perfstat\_cpu\_total**, **perfstat\_disk**, **perfstat\_diskadapter**, **perfstat\_netinterface**, and **perfstat\_pagingspace** subroutines return configuration information retrieved from the ODM database and automatically cached by the library.

The **perfstat\_reset** subroutine flushes this information cache and should be called whenever the machine configuration has changed.

## Files

The **libperfstat.h** defines standard macros, data types and subroutines.

## Related Information

“perfstat\_cpu\_total Subroutine” on page 1004, “perfstat\_disk Subroutine” on page 1006, “perfstat\_diskadapter Subroutine” on page 1008, “perfstat\_diskpath Subroutine” on page 1009, “perfstat\_netinterface Subroutine” on page 1015, “perfstat\_pagingspace Subroutine” on page 1018, and “perfstat\_partial\_reset Subroutine” on page 1019.

Perfstat API in Performance Tools and APIs Technical Reference.

---

## perror Subroutine

### Purpose

Writes a message explaining a subroutine error.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <errno.h>
#include <stdio.h>

void perror ( String)
const char *String;
extern int errno;
extern char *sys_errlist[ ];
extern int sys_nerr;
```

### Description

The **perror** subroutine writes a message on the standard error output that describes the last error encountered by a system call or library subroutine. The error message includes the *String* parameter string followed by a : (colon), a space character, the message, and a new-line character. The *String* parameter



string should include the name of the program that caused the error. The error number is taken from the **errno** global variable, which is set when an error occurs but is not cleared when a successful call to the  **perror** subroutine is made.

To simplify various message formats, an array of message strings is provided in the **sys\_errlist** structure or use the **errno** global variable as an index into the **sys\_errlist** structure to get the message string without the new-line character. The largest message number provided in the table is **sys\_nerr**. Be sure to check the **sys\_nerr** structure because new error codes can be added to the system before they are added to the table.

The  **perror** subroutine retrieves an error message based on the language of the current locale.

After successfully completing, and before a call to the **exit** or **abort** subroutine or the completion of the **fflush** or **fclose** subroutine on the standard error stream, the  **perror** subroutine marks for update the **st\_ctime** and **st\_mtime** fields of the file associated with the standard error stream.

## Parameter

*String* Specifies a parameter string that contains the name of the program that caused the error. The ensuing printed message contains this string, a : (colon), and an explanation of the error.

## Related Information

The **abort** subroutine, **exit** subroutine, **fflush** or **fclose** subroutine, **printf**, **fprintf**, **sprintf**, **wsprintf**, **vprintf**, **vfprintf**, **vsprintf**, or **vwsprintf** subroutine, **strerror** subroutine.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pipe Subroutine

### Purpose

Creates an interprocess channel.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
```

```
int pipe ( FileDescriptor )
int FileDescriptor[2];
```

### Description

The **pipe** subroutine creates an interprocess channel called a pipe and returns two file descriptors, *FileDescriptor*[0] and *FileDescriptor*[1]. *FileDescriptor*[0] is opened for reading and *FileDescriptor*[1] is opened for writing.

A read operation on the *FileDescriptor*[0] parameter accesses the data written to the *FileDescriptor*[1] parameter on a first-in, first-out (FIFO) basis.

Write requests of **PIPE\_BUF** bytes or fewer will not be interleaved (mixed) with data from other processes doing writes on the same pipe. **PIPE\_BUF** is a system variable described in the **pathconf** ("pathconf or

fpathconf Subroutine” on page 979) subroutine. Writes of greater than **PIPE\_BUF** bytes may have data interleaved, on arbitrary boundaries, with other writes.

If **O\_NONBLOCK** or **O\_NDELAY** are set, writes requests of **PIPE\_BUF** bytes or fewer will either succeed completely or fail and return -1 with the **errno** global variable set to **EAGAIN**. A write request for more than **PIPE\_BUF** bytes will either transfer what it can and return the number of bytes actually written, or transfer no data and return -1 with the **errno** global variable set to **EAGAIN**.

## Parameters

*FileDescriptor* Specifies the address of an array of two integers into which the new file descriptors are placed.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned, and the **errno** global variable is set to identify the error.

## Error Codes

The **pipe** subroutine is unsuccessful if one or more the following are true:

**EFAULT** The *FileDescriptor* parameter points to a location outside of the allocated address space of the process.  
**EMFILE** The number of open of file descriptors exceeds the **OPEN\_MAX** value.  
**ENFILE** The system file table is full, or the device containing pipes has no free i-nodes.

## Related Information

The **read** subroutine, **select** subroutine, **write** subroutine.

The **ksh** command, **sh** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## plock Subroutine

### Purpose

Locks the process, text, or data in memory.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/lock.h>
```

```
int plock ( Operation )  
int Operation;
```

### Description

The **plock** subroutine allows the calling process to lock or unlock its text region (text lock), its data region (data lock), or both its text and data regions (process lock) into memory. The **plock** subroutine does not lock the shared text segment or any shared data segments. Locked segments are pinned in memory and

are immune to all routine paging. Memory locked by a parent process is not inherited by the children after a **fork** subroutine call. Likewise, locked memory is unlocked if a process executes one of the **exec** subroutines. The calling process must have the root user authority to use this subroutine.

A real-time process can use this subroutine to ensure that its code, data, and stack are always resident in memory.

**Note:** Before calling the **plock** subroutine, the user application must lower the maximum stack limit value using the **ulimit** subroutine.

## Parameters

<i>Operation</i>	Specifies one of the following:
	<b>PROCLOCK</b> Locks text and data into memory (process lock).
	<b>TXTLOCK</b> Locks text into memory (text lock).
	<b>DATLOCK</b> Locks data into memory (data lock).
	<b>UNLOCK</b> Removes locks.

## Return Values

Upon successful completion, a value of 0 is returned to the calling process. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **plock** subroutine is unsuccessful if one or more of the following is true:

<b>EPERM</b>	The effective user ID of the calling process does not have the root user authority.
<b>EINVAL</b>	The <i>Operation</i> parameter has a value other than <b>PROCLOCK</b> , <b>TXTLOCK</b> , <b>DATLOCK</b> , or <b>UNLOCK</b> .
<b>EINVAL</b>	The <i>Operation</i> parameter is equal to <b>PROCLOCK</b> , and a process lock, text lock, or data lock already exists on the calling process.
<b>EINVAL</b>	The <i>Operation</i> parameter is equal to <b>TXTLOCK</b> , and a text lock or process lock already exists on the calling process.
<b>EINVAL</b>	The <i>Operation</i> parameter is equal to <b>DATLOCK</b> , and a data lock or process lock already exists on the calling process.
<b>EINVAL</b>	The <i>Operation</i> parameter is equal to <b>UNLOCK</b> , and no type of lock exists on the calling process.

## Related Information

The **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236) subroutines, **\_exit**, **exit**, or **atexit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **ulimit** subroutine.

---

## pm\_cycles Subroutine

### Purpose

Measures processor speed in cycles per second.

## Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>
```

```
double pm_cycles (void)
```

## Description

The **pm\_cycles** subroutine uses the Performance Monitor cycle counter and the processor real-time clock to measure the actual processor clock speed. The speed is returned in cycles per second.

## Return Values

0	An error occurred.
Processor speed in cycles per second	No errors occurred.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_delete\_program Subroutine

### Purpose

Deletes previously established systemwide Performance Monitor settings.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program ()
```

### Description

The **pm\_delete\_program** subroutine deletes previously established systemwide Performance Monitor settings.

### Return Values

0	No errors occurred.
Positive error code	Refer to the <b>pm_error</b> ("pm_error Subroutine" on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

`pm_init` (“pm\_init Subroutine” on page 1084), `pm_error` (“pm\_error Subroutine” on page 1035), `pm_set_program` (“pm\_set\_program Subroutine” on page 1095), `pm_get_program` (“pm\_get\_program Subroutine” on page 1061), `pm_get_data` (“pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine” on page 1036), `pm_start` (“pm\_start and pm\_tstart Subroutine” on page 1121), `pm_stop` (“pm\_stop and pm\_tstop Subroutine ” on page 1130), `pm_reset_data` (“pm\_reset\_data Subroutine” on page 1087) subroutines.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_delete\_program\_group Subroutine

### Purpose

Deletes previously established Performance Monitor settings for the counting group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program_group ( pid, tid)
pid_t pid;
tid_t tid;
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_delete\_program\_pgroup** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_delete\_program\_pgroup** subroutine with a *ptid* parameter equal to 0.

The **pm\_delete\_program\_group** subroutine deletes previously established Performance Monitor settings for a target kernel thread. The thread must be stopped and must be part of a debuggee process under the control of the calling process. The settings for the group to which the target thread belongs and from all the other threads in the same group are also deleted.

### Parameters

*pid*

Process identifier of target thread. The target process must be a debuggee under the control of the calling process.

*tid*

Thread identifier of a target thread.

## Return Values

**0** No errors occurred.  
**Positive error code** Refer to the `pm_error` (“`pm_error Subroutine`” on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the `pm_error` (“`pm_error Subroutine`” on page 1035) subroutine.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The `pm_init` (“`pm_init Subroutine`” on page 1084) subroutine, `pm_error` (“`pm_error Subroutine`” on page 1035) subroutine, `pm_set_program_group` (“`pm_set_program_group Subroutine`” on page 1096) subroutine, `pm_get_program_group` (“`pm_get_program_group Subroutine`” on page 1062) subroutine, `pm_get_data_group` (“`pm_get_data_group, pm_get_tdata_group and pm_get_Tdata_group Subroutine`” on page 1039) subroutine, `pm_start_group` (“`pm_start_group and pm_tstart_group Subroutine`” on page 1122) subroutine, `pm_stop_group` (“`pm_stop_group and pm_tstop_group Subroutine`” on page 1131) subroutine, `pm_reset_data_group` (“`pm_reset_data_group Subroutine`” on page 1088) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## `pm_delete_program_mygroup` Subroutine

### Purpose

Deletes previously established Performance Monitor settings for the counting group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (**`libpmapi.a`**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program_mygroup ()
```

### Description

The `pm_delete_program_mygroup` subroutine deletes previously established Performance Monitor settings for the calling kernel thread, the counting group to which it belongs, and for all the threads that are members of the same group.

## Return Values

**0** No errors occurred.

**Positive error code** Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

`pm_init` (“pm\_init Subroutine” on page 1084), `pm_error` (“pm\_error Subroutine” on page 1035), `pm_set_program_mygroup` (“pm\_set\_program\_mygroup Subroutine” on page 1102), `pm_get_program_mygroup` (“pm\_get\_program\_mygroup Subroutine” on page 1068), `pm_get_data_mygroup` (“pm\_get\_data\_mygroup, pm\_get\_tdata\_mygroup or pm\_get\_Tdata\_mygroup Subroutine” on page 1044), `pm_start_mygroup` (“pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123), `pm_stop_mygroup` (“pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132), `pm_reset_data_mygroup` (“pm\_reset\_data\_mygroup Subroutine” on page 1090) subroutines.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_delete\_program\_mythread Subroutine

### Purpose

Deletes the previously established Performance Monitor settings for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program_mythread ()
```

### Description

The **pm\_delete\_program\_mythread** subroutine deletes the previously established Performance Monitor settings for the calling kernel thread.

### Return Values

0 No errors occurred.

**Positive error code** Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

`pm_init` (“`pm_init` Subroutine” on page 1084), `pm_error` (“`pm_error` Subroutine” on page 1035), `pm_set_program_mythread` (“`pm_set_program_mythread` Subroutine” on page 1106), `pm_get_program_mythread` (“`pm_get_program_mythread` Subroutine” on page 1071), `pm_get_data_mythread` (“`pm_get_data_mythread`, `pm_get_tdata_mythread` or `pm_get_Tdata_mythread` Subroutine” on page 1047), `pm_start_mythread` (“`pm_start_mythread` and `pm_tstart_mythread` Subroutine” on page 1125), `pm_stop_mythread` (“`pm_stop_mythread` and `pm_tstop_mythread` Subroutine ” on page 1133), `pm_reset_data_mythread` (“`pm_reset_data_mythread` Subroutine” on page 1090) subroutines.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## `pm_delete_program_pgroup` Subroutine

### Purpose

Deletes previously established Performance Monitor settings for the counting group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (`libpmapi.a`)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program_pgroup ( pid, tid, ptid)
pid_t pid;
tid_t tid;
ptid_t ptid;
```

### Description

The `pm_delete_program_pgroup` subroutine deletes previously established Performance Monitor settings for a target pthread. The pthread must be stopped and must be part of a debuggee process under the control of the calling process. The settings for the group to which the target pthread belongs and from all the other pthreads in the same group are also deleted.

If the pthread is running in 1:1 mode, only the `tid` parameter must be specified. If the pthread is running in m:n mode, only the `ptid` parameter must be specified. If both the `ptid` and `tid` parameters are specified, they must be referring to a single pthread with the `ptid` parameter specified and currently running on a kernel thread with specified `tid` parameter.

### Parameters

<code>pid</code>	Process ID of target thread. The target process must be a debuggee under the control of the calling process.
<code>tid</code>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<code>ptid</code>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.



## Return Values

0 No errors occurred.  
Positive error code Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_set\_program\_pgroup Subroutine” on page 1109, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, and the “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_delete\_program\_pthread Subroutine

### Purpose

Deletes the previously established Performance Monitor settings for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program_pthread ( pid, tid, ptid)
pid_t pid;
tid_t tid;
ptid_t ptid;
```

### Description

The **pm\_delete\_program\_pthread** subroutine deletes the previously established Performance Monitor settings for a target pthread. The pthread must be stopped and must be part of a debuggee process under the control of the calling process.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee under the control of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_delete\_program\_pthread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pthread, pm\_get\_tdata\_pthread or pm\_get\_Tdata\_pthread Subroutine” on page 1054, “pm\_get\_program\_pthread Subroutine” on page 1077, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pthread Subroutine” on page 1093, “pm\_set\_program\_pthread Subroutine” on page 1113, “pm\_start\_pthread and pm\_tstart\_pthread Subroutine” on page 1127, “pm\_stop\_pthread and pm\_tstop\_pthread Subroutine ” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_delete\_program\_thread Subroutine

### Purpose

Deletes the previously established Performance Monitor settings for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_delete_program_thread ( pid, tid)
pid_t pid;
tid_t tid;
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_delete\_program\_pthread** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_delete\_program\_pthread** subroutine with a *ptid* parameter equal to 0.

The `pm_delete_program_thread` subroutine deletes the previously established Performance Monitor settings for a target kernel thread. The thread must be stopped and must be part of a debuggee process under the control of the calling process.

## Parameters

<code>pid</code>	Process identifier of target thread. Target process must be a debuggee under the control of the calling process.
<code>tid</code>	Thread identifier of the target thread.

## Return Values

<code>0</code>	No errors occurred.
<b>Positive error code</b>	Refer to the <code>pm_error</code> (“ <code>pm_error</code> Subroutine”) subroutine to decode the error code.

## Error Codes

Refer to the `pm_error` (“`pm_error` Subroutine”) subroutine.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

`pm_init` (“`pm_init` Subroutine” on page 1084), `pm_error` (“`pm_error` Subroutine”), `pm_set_program_thread` (“`pm_set_program_thread` Subroutine” on page 1117), `pm_get_program_thread` (“`pm_get_program_thread` Subroutine” on page 1080), `pm_get_data_thread` (“`pm_get_data_thread`, `pm_get_tdata_thread` or `pm_get_Tdata_thread` Subroutine” on page 1057), `pm_start_thread` (“`pm_start_thread` and `pm_tstart_thread` Subroutine” on page 1129), `pm_stop_thread` (“`pm_stop_thread` and `pm_tstop_thread` Subroutine” on page 1137), `pm_reset_data_thread` (“`pm_reset_data_thread` Subroutine” on page 1094) subroutines.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_error Subroutine

### Purpose

Decodes Performance Monitor APIs error codes.

### Library

Performance Monitor APIs Library (`libpmapi.a`)

### Syntax

```
#include <pmapi.h>
```

```
void pm_error ( *Where, errorcode)
char *Where;
int errorcode;
```

## Description

The **pm\_error** subroutine writes a message on the standard error output that describes the parameter *errorcode* encountered by a Performance Monitor API library subroutine. The error message includes the *Where* parameter string followed by a : (colon), a space character, the message, and a new-line character. The *Where* parameter string includes the name of the program that caused the error.

## Parameters

<i>*Where</i>	Specifies where the error was encountered.
<i>errorcode</i>	Specifies the error code as returned by one of the Performance Monitor APIs library subroutines.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The **pm\_init** subroutine, **pm\_set\_program** subroutine, **pm\_get\_program** subroutine, **pm\_delete\_program** subroutine, **pm\_get\_data** subroutine, **pm\_start** subroutine, **pm\_stop** subroutine, **pm\_reset\_data** subroutine.

The **pm\_set\_program\_mythread** subroutine, **pm\_get\_program\_mythread** subroutine, **pm\_delete\_program\_mythread** subroutine, **pm\_get\_data\_mythread** subroutine, **pm\_start\_mythread** subroutine, **pm\_stop\_mythread** subroutine, **pm\_reset\_data\_mythread** subroutine.

The **pm\_set\_program\_mygroup** subroutine, **pm\_get\_program\_mygroup** subroutine, **pm\_delete\_program\_mygroup** subroutine, **pm\_get\_data\_mygroup** subroutine, **pm\_start\_mygroup** subroutine, **pm\_stop\_mygroup** subroutine, **pm\_reset\_data\_mygroup** subroutine.

The **pm\_set\_program\_thread** subroutine, **pm\_get\_program\_thread** subroutine, **pm\_delete\_program\_thread** subroutine, **pm\_get\_data\_thread** subroutine, **pm\_start\_thread** subroutine, **pm\_stop\_thread** subroutine, **pm\_reset\_data\_thread** subroutine.

The **pm\_set\_program\_group** subroutine, **pm\_get\_program\_group** subroutine, **pm\_delete\_program\_group** subroutine, **pm\_get\_data\_group** subroutine, **pm\_start\_group** subroutine, **pm\_stop\_group** subroutine, **pm\_reset\_data\_group** subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## **pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine**

### Purpose

Returns systemwide Performance Monitor data.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>

int pm_get_data ( *pmdata)
pm_data_t *pmdata;

int pm_get_tdata (pmdata, * time)
pm_data_t *pmdata;
timebasestruct_t *time;

int pm_get_Tdata (pmdata, * times)
pm_data_t *pmdata;
pm_accu_time_t *times;

int pm_get_data_cpu (cpuid, *pmdata)
int cpuid;
pm_data_t *pmdata;

int pm_get_tdata_cpu (cpuid, *pmdata, *time)
int cpuid;
pm_data_t *pmdata;
timebasestruct_t *time;

int pm_get_Tdata_cpu (cpuid, *pmdata, *times)
int cpuid;
pm_data_t *pmdata;
pm_accu_time_t *times
int pm_get_data_lcpu (lcpuid, *pmdata)
int lcpuid;
pm_data_t *pmdata;

int pm_get_tdata_lcpu (lcpuid, *pmdata, *time)
int lcpuid;
pm_data_t *pmdata;
timebasestruct_t *time;

int pm_get_Tdata_lcpu (lcpuid, *pmdata, *times)
int lcpuid;
pm_data_t *pmdata;
pm_accu_time_t *times
```

## Description

The **pm\_get\_data** subroutine retrieves the current systemwide Performance Monitor data.

The **pm\_get\_tdata** subroutine retrieves the current systemwide Performance Monitor data, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata** subroutine retrieves the current systemwide Performance Monitor data, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The **pm\_get\_data\_cpu**, **pm\_get\_tdata\_cpu**, and **pm\_get\_Tdata\_cpu** subroutines retrieve the current Performance Monitor data for a specified processor. The given processor ID represents a contiguous number ranging from 0 to `_system_configuration.ncpus`. These subroutines can only be used when no Dynamic Reconfiguration operations are made on the machine, because when processors are added or removed, the processor numbering is modified and the specified processor number can designate different processors from one call to another. These subroutines are maintained for compatibility with previous versions.

The **pm\_get\_data\_cpu** subroutine retrieves the current Performance Monitor data for the specified processor.

The `pm_get_tdata_cpu` subroutine retrieves the current Performance Monitor data for the specified processor, and a timestamp indicating the last time the hardware counters were read.

The `pm_get_Tdata_cpu` subroutine retrieves the current Performance Monitor data for the specified processor, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The `pm_get_data_lcpu`, `pm_get_tdata_lcpu`, and `pm_get_Tdata_lcpu` subroutines retrieve the current Performance Monitor data for a specified logical processor. The given processor ID represents a value ranging from 0 to `_system_configuration.max_ncpus`. This value always represents the same processor, even after Dynamic Reconfiguration operations have occurred. These subroutines might return an error if the specified logical processor number has never run during the counting interval.

The `pm_get_data_lcpu` subroutine retrieves the current Performance Monitor data for the specified logical processor.

The `pm_get_tdata_lcpu` subroutine retrieves the current Performance Monitor data for the specified logical processor, and a timestamp indicating the last time the hardware counters were read.

The `pm_get_Tdata_lcpu` subroutine retrieves the current Performance Monitor data for the specified logical processor, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The Performance Monitor data is always a set (one per hardware counter on the machines used) of 64-bit values.

## Parameters

<i>*pmdata</i>	Pointer to a structure that contains the returned systemwide Performance Monitor data.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <code>time_base_to_time</code> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <code>time_base_to_time</code> subroutine.
<i>cpuid</i>	Contiguous processor numbers ranging from 0 to <code>_system_configuration.ncpus</code> . This value does not always designate the same processor, even after Dynamic Reconfiguration operations have occurred.
<i>lcpuid</i>	Logical processor identifier. Each identifier stays linked to a particular processor between reboots, even after Dynamic Reconfiguration operations. This value must be in the range from 0 to <code>_system_configuration.max_ncpus</code> .

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the <code>pm_error</code> Subroutine to decode the error code.

## Error Codes

Refer to the `pm_error` Subroutine.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program Subroutine” on page 1095, “pm\_get\_program Subroutine” on page 1061, “pm\_delete\_program Subroutine” on page 1028, “pm\_start and pm\_tstart Subroutine” on page 1121, “pm\_stop and pm\_tstop Subroutine ” on page 1130, “pm\_reset\_data Subroutine” on page 1087.

The **read\_real\_time or time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine

### Purpose

Returns Performance Monitor data for the counting group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_group (pid, tid, *pmdata)
pid_t pid;
tid_t tid;
pm_data_t *pmdata;
```

```
int pm_get_tdata_group (pid, tid, *pmdata, *time)
pm_data_t *pmdata;
pid_t pid;
tid_t tid;
timebasestruct_t *time;
```

```
int pm_get_Tdata_group (pid, tid, *pmdata, *times)
pm_data_t *pmdata;
pid_t pid;
tid_t tid;
pm_accu_time_t *times;
```

### Description

These subroutines support only the 1:1 threading model. They have been superseded by the **pm\_get\_data\_pgroup** and **pm\_get\_tdata\_pgroup** subroutines, which support both the 1:1 and the M:N threading models. Calls to these subroutines are equivalent to calls to the **pm\_get\_data\_pgroup** and **pm\_get\_tdata\_pgroup** subroutines with a *ptid* parameter equal to 0.

The **pm\_get\_data\_group** subroutine retrieves the current Performance Monitor data for the counting group to which a target kernel thread belongs. The thread must be stopped and must be part of a debuggee process under the control of the calling process.

The **pm\_get\_tdata\_group** subroutine retrieves the current Performance Monitor data for the counting group to which a target thread belongs, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata\_group** subroutine retrieves the current Performance Monitor data for the counting group to which a target thread belongs, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The Performance Monitor data is always a set (one per hardware counter on the machine used) of 64-bit values. The information returned also includes the characteristics of the group, such as the number of its members, if it is a process level group, and if its counters are consistent with the sum of the counters for all of the threads in the group.

## Parameters

<i>pid</i>	Process identifier of a target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread identifier of a target thread.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data for the group to which the target thread belongs.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_group Subroutine” on page 1096, “pm\_get\_program\_group Subroutine” on page 1062, “pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine” on page 1039, “pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122, “pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131, “pm\_reset\_data\_group Subroutine” on page 1088.

The **read\_real\_time** or **time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.



---

## pm\_get\_data\_group\_mx and pm\_get\_tdata\_group\_mx Subroutine

### Purpose

Returns Performance Monitor data in counter multiplexing mode for the counting group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_group_mx (pid, tid, *pmdata)
pid_t pid;
tid_t tid;
pm_data_mx_t *pmdata;
```

```
int pm_get_tdata_group_mx (pid, tid, *pmdata, *time)
pm_data_mx_t *pmdata;
pid_t pid;
tid_t tid;
timebasestruct_t *time;
```

### Description

These subroutines support only the 1:1 threading model. They have been superseded by the **pm\_get\_data\_pgroup\_mx** and **pm\_get\_tdata\_pgroup\_mx** subroutines, which support both the 1:1 and the M:N threading models. Calls to these subroutines are equivalent to calls to the **pm\_get\_data\_pgroup\_mx** and **pm\_get\_tdata\_pgroup\_mx** subroutines with a *ptid* parameter equal to 0.

The **pm\_get\_data\_group\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the counting group to which a target kernel thread belongs. The thread must be stopped and must be part of a debuggee process under the control of the calling process.

The **pm\_get\_tdata\_group\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the counting group to which a target thread belongs, and a timestamp indicating the last time the hardware counters were read.

The Performance Monitor data is always an array of a set (one per hardware counter on the machine used) of 64-bit values. The information returned also includes the characteristics of the group, such as the number of its members, whether it is a process level group, and whether its counters are consistent with the sum of the counters for all of the threads in the group.

The user application must free the array allocated to store accumulated counts and times (the *accu\_set* field of the *pmdata* parameter).

### Parameters

<i>pid</i>	Process identifier of a target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread identifier of a target thread.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted) for the group to which the target thread belongs.

*\*time*

Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

**0** No errors occurred.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_group\_mx and pm\_set\_program\_group\_mm Subroutines” on page 1098, “pm\_get\_program\_group\_mx and pm\_get\_program\_group\_mm Subroutines” on page 1064, “pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122, “pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131, and “pm\_reset\_data\_group Subroutine” on page 1088.

The **read\_real\_time** or **time\_base\_to\_time** subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_mx, pm\_get\_tdata\_mx, pm\_get\_data\_cpu\_mx, pm\_get\_tdata\_cpu\_mx, pm\_get\_data\_lcpu\_mx and pm\_get\_tdata\_lcpu\_mx Subroutine

### Purpose

Returns systemwide Performance Monitor data in counter multiplexing mode.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_mx ( *pdata)  
pm_data_mx_t *pdata;
```

```
int pm_get_tdata_mx (pdata, *time)  
pm_data_mx_t *pdata;  
timebasestruct_t *time;
```

```
int pm_get_data_cpu_mx (cpuid, *pdata)  
int cpuid;
```

```

pm_data_mx_t *pmdata;

int pm_get_tdata_cpu_mx (cpuid, *pmdata, *time)
int cpuid;
pm_data_mx_t *pmdata;
timebasestruct_t *time;

int pm_get_data_lcpu_mx (lcpuid, *pmdata)
int lcpuid;
pm_data_mx_t *pmdata;

int pm_get_tdata_lcpu_mx (lcpuid, *pmdata, *time)
int lcpuid;
pm_data_mx_t *pmdata;
timebasestruct_t *time;

```

## Description

The **pm\_get\_data\_mx** subroutine retrieves the current systemwide Performance Monitor data in counter multiplexing mode.

The **pm\_get\_tdata\_mx** subroutine retrieves the current systemwide Performance Monitor data in counter multiplexing mode, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_data\_cpu\_mx** and the **pm\_get\_tdata\_cpu\_mx** subroutines retrieve the current Performance Monitor data for a specified processor. The given processor ID represents a contiguous number ranging from 0 to `_system_configuration.ncpus`. These subroutines can only be used when no Dynamic Reconfiguration operations are made on the machine, because when processors are added or removed, the processor numbering is modified and the specified processor number can designate different processors from one call to another. These subroutines are maintained for compatibility with previous versions.

The **pm\_get\_data\_cpu\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the specified processor.

The **pm\_get\_tdata\_cpu\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the specified processor, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_data\_lcpu\_mx** and the **pm\_get\_tdata\_lcpu\_mx** subroutines retrieve the current Performance Monitor data for a specified logical processor. The given processor ID represents a value ranging from 0 to `_system_configuration.max_ncpus`. This value always represents the same processor, even after Dynamic Reconfiguration operations have occurred. These subroutines might return an error if the specified logical processor number has never run during the counting interval.

The **pm\_get\_data\_lcpu\_mx** subroutine retrieves the current Performance Monitor data for the specified logical processor in counter multiplexing mode.

The **pm\_get\_tdata\_lcpu\_mx** subroutine retrieves the current Performance Monitor data for the specified logical processor in counter multiplexing mode, and a timestamp indicating the last time the hardware counters were read.

The Performance Monitor data is always an array of a set (one per hardware counter on the machines used) of 64-bit values.

The user application must free the array allocated to store accumulated counts and times (the `accu_set` field of the `pmdata` parameter).

## Parameters

<i>*pdata</i>	Pointer to a structure that contains the returned systemwide Performance Monitor data. (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted)
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>cpuid</i>	Contiguous processor numbers going from 0 to <code>_system_configuration.ncpus</code> . This value does not always designate the same processor, even after Dynamic Reconfiguration operations have occurred.
<i>lcpuid</i>	Logical processor identifier. Each identifier stays linked to a particular processor between reboots, even after Dynamic Reconfiguration operations. This value must be in the range from 0 to <code>_system_configuration.max_ncpus</code> .

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the <code>pm_error</code> Subroutine to decode the error code.

## Error Codes

Refer to the `pm_error` Subroutine.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “`pm_init` Subroutine” on page 1084, “`pm_error` Subroutine” on page 1035, “`pm_set_program_mx` and `pm_set_program_mm` Subroutines” on page 1100, “`pm_get_program_mx` and `pm_get_program_mm` Subroutines” on page 1066, “`pm_delete_program` Subroutine” on page 1028, “`pm_start` and `pm_tstart` Subroutine” on page 1121, “`pm_stop` and `pm_tstop` Subroutine ” on page 1130, and the “`pm_reset_data` Subroutine” on page 1087.

The **`read_real_time` or `time_base_to_time`** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## `pm_get_data_mygroup`, `pm_get_tdata_mygroup` or `pm_get_Tdata_mygroup` Subroutine

### Purpose

Returns Performance Monitor data for the counting group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (**`libpmapi.a`**)

## Syntax

```
#include <pmapi.h>

int pm_get_data_mygroup (*pmdata)
pm_data_t *pmdata;

int pm_get_tdata_mygroup (*pmdata, *time)
pm_data_t *pmdata;
timebasestruct_t *time;

int pm_get_Tdata_mygroup (pmdata, *times)
pm_data_t *pmdata;
pm_accu_time_t *times;
```

## Description

The **pm\_get\_data\_mygroup** subroutine retrieves the current Performance Monitor data for the group to which the calling kernel thread belongs.

The **pm\_get\_tdata\_mygroup** subroutine retrieves the current Performance Monitor data for the group to which the calling thread belongs, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata\_mygroup** subroutine retrieves the current Performance Monitor data for the group to which the calling thread belongs, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The Performance Monitor data is always a set (one per hardware counter on the machine used) of 64-bit values. The information returned also includes the characteristics of the group, such as the number of its members, if it is a process level group, and if its counters are consistent with the sum of the counters for all of the threads in the group.

## Parameters

<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data for the group to which the calling thread belongs.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mygroup Subroutine” on page 1102, “pm\_get\_program\_mygroup Subroutine” on page 1068, “pm\_get\_data\_mygroup, pm\_get\_tdata\_mygroup or pm\_get\_Tdata\_mygroup Subroutine” on page 1044, “pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123, “pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132, “pm\_reset\_data\_mygroup Subroutine” on page 1090.

The **read\_real\_time or time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_mygroup\_mx or pm\_get\_tdata\_mygroup\_mx Subroutine

### Purpose

Returns Performance Monitor data in counter multiplexing mode for the counting group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_mygroup_mx (*pmdata)
pm_data_mx_t *pmdata;
```

```
int pm_get_tdata_mygroup_mx (*pmdata, *time)
pm_data_mx_t *pmdata;
timebasestruct_t *time;
```

### Description

The **pm\_get\_data\_mygroup\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the group to which the calling kernel thread belongs.

The **pm\_get\_tdata\_mygroup\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the group to which the calling thread belongs, and a timestamp indicating the last time the hardware counters were read.

The Performance Monitor data is always an array of set (one per hardware counter on the machine used) of 64-bit values. The information returned also includes the characteristics of the group, such as the number of its members, if it is a process level group, and if its counters are consistent with the sum of the counters for all of the threads in the group.

The user application must free the array allocated to store accumulated counts and times (accu\_set field of pmdata).

## Parameters

<i>*pdata</i>	Pointer to a structure to return the Performance Monitor data (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted) for the group to which the calling thread belongs.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mygroup\_mx and pm\_set\_program\_mygroup\_mm Subroutines” on page 1103, “pm\_get\_program\_mygroup\_mx and pm\_get\_program\_mygroup\_mm Subroutines” on page 1069, “pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123, “pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132, “pm\_reset\_data\_mygroup Subroutine” on page 1090.

The **read\_real\_time** or **time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine

### Purpose

Returns Performance Monitor data for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_mythread (*pdata)  
pm_data_t *pdata;
```

```
int pm_get_tdata_mythread (*pmdata, *time)
pm_data_t *pmdata;
timebasestruct_t *time;
```

```
int pm_get_Tdata_mythread (pmdata, * times)
pm_data_t *pmdata;
pm_accu_time_t *times;
```

## Description

The **pm\_get\_data\_mythread** subroutine retrieves the current Performance Monitor data for the calling kernel thread.

The **pm\_get\_tdata\_mythread** subroutine retrieves the current Performance Monitor data for the calling kernel thread, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata\_mythread** subroutine retrieves the current Performance Monitor data for the calling kernel thread, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The Performance Monitor data is always a set (one per hardware counter on the machine used) of 64-bit values.

## Parameters

<i>*pmdata</i>	Pointer to a structure to contain the returned Performance Monitor data for the calling kernel thread.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<b>/usr/include/pmapi.h</b>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mythread Subroutine” on page 1106, “pm\_get\_program\_mythread Subroutine” on page 1071, “pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine” on page 1047), “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133, “pm\_reset\_data\_mythread Subroutine” on page 1090.



The `read_real_time` or `time_base_to_time` subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_mythread\_mx or pm\_get\_tdata\_mythread\_mx Subroutine

### Purpose

Returns Performance Monitor data in counter multiplexing mode for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_get_data_mythread_mx (*pmdata)
pm_data_mx_t *pmdata;

int pm_get_tdata_mythread_mx (*pmdata, *time)
pm_data_mx_t *pmdata;
timebasestruct_t *time;
```

### Description

The `pm_get_data_mythread_mx` subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the calling kernel thread.

The `pm_get_tdata_mythread_mx` subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the calling kernel thread, and a timestamp indicating the last time the hardware counters were read.

The Performance Monitor data is always an array of a set (one per hardware counter on the machine used) of 64-bit values.

The user application must free the array allocated to store accumulated counts and times (the `accu_set` field of the `pmdata` parameter).

### Parameters

<i>*pmdata</i>	Pointer to a structure to contain the returned Performance Monitor data (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted) for the calling kernel thread.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <code>time_base_to_time</code> subroutine.

### Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mythread\_mx and pm\_set\_program\_mythread\_mm Subroutines” on page 1107, “pm\_get\_program\_mythread\_mx and pm\_get\_program\_mythread\_mm Subroutines” on page 1072, “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133, “pm\_reset\_data\_mythread Subroutine” on page 1090.

The **read\_real\_time\_or\_time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine

### Purpose

Returns Performance Monitor data for the counting group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_pgroup (pid, tid, ptid, *pmdata)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_t *pmdata;
```

```
int pm_get_tdata_pgroup (pid, tid, *pmdata, *time)
pm_data_t *pmdata;
pid_t pid;
tid_t tid;
ptid_t ptid;
timebasestruct_t *time;
```

```
int pm_get_Tdata_pgroup (pid, tid, *pmdata, *times)
pm_data_t *pmdata;
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_accu_time_t *times;
```

## Description

The **pm\_get\_data\_pgroup** subroutine retrieves the current Performance Monitor data for the counting group to which a target pthread belongs. The pthread must be stopped and must be part of a debuggee process under the control of the calling process.

The **pm\_get\_tdata\_pgroup** subroutine retrieves the current Performance Monitor data for the counting group to which a target pthread belongs, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata\_pgroup** subroutine retrieves the current Performance Monitor data for the counting group to which a target pthread belongs, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The Performance Monitor data is always a set (one per hardware counter on the machine used) of 64-bit values. The information returned also includes the characteristics of the group, such as the number of its members, if it is a process level group, and if its counters are consistent with the sum of the counters for all of the pthreads in the group.

## Parameters

<i>pid</i>	Process identifier of a target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data for the group to which the target pthread belongs.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_delete\_program\_pthread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_get\_program\_pgroup Subroutine” on page 1074, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_set\_program\_pgroup Subroutine” on page 1109, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

The **read\_real\_time or time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_pgroup\_mx and pm\_get\_tdata\_pgroup\_mx Subroutine

### Purpose

Returns Performance Monitor data in counter multiplexing mode for the counting group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_pgroup_mx (pid, tid, ptid, *pdata)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_mx_t *pdata;
```

```
int pm_get_tdata_pgroup_mx (pid, tid, *pdata, *time)
pm_data_mx_t *pdata;
pid_t pid;
tid_t tid;
ptid_t ptid;
timebasestruct_t *time;
```

### Description

The **pm\_get\_data\_pgroup\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the counting group to which a target pthread belongs. The pthread must be stopped and must be part of a debuggee process under the control of the calling process.

The **pm\_get\_tdata\_pgroup\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for the counting group to which a target pthread belongs, and a timestamp indicating the last time the hardware counters were read.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The Performance Monitor data is always an array of a set (one per hardware counter on the machine used) of 64-bit values. The information returned also includes the characteristics of the group, such as the

number of its members, whether it is a process level group, and whether its counters are consistent with the sum of the counters for all of the pthreads in the group.

The user application must free the array allocated to store accumulated counts and times (the `accu_set` field of the `pmdata` parameter).

## Parameters

<code>pid</code>	Process identifier of a target thread. The target process must be an argument of a debug process.
<code>tid</code>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<code>ptid</code>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<code>*pmdata</code>	Pointer to a structure to return the Performance Monitor data (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted) for the group to which the target pthread belongs.
<code>*time</code>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <code>time_base_to_time</code> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “ <code>pm_error</code> Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “`pm_error` Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “`pm_delete_program_pthread` Subroutine” on page 1033, “`pm_error` Subroutine” on page 1035, “`pm_get_program_pgroup_mx` and `pm_get_program_pgroup_mm` Subroutines” on page 1075, “`pm_initialize` Subroutine” on page 1085, “`pm_reset_data_pgroup` Subroutine” on page 1091, “`pm_set_program_pgroup_mx` and `pm_set_program_pgroup_mm` Subroutines” on page 1111, “`pm_start_pgroup` and `pm_tstart_pgroup` Subroutine” on page 1126, “`pm_stop_pgroup` and `pm_tstop_pgroup` Subroutine” on page 1134.

The `read_real_time` or `time_base_to_time` subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_pthread, pm\_get\_tdata\_pthread or pm\_get\_Tdata\_pthread Subroutine

### Purpose

Returns Performance Monitor data for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_pthread (pid, tid, ptid, *pdata)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_t *pdata;
```

```
int pm_get_tdata_pthread (pid, tid, ptid, *pdata, *time)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_t *pdata;
timebasestruct_t *time;
```

```
int pm_get_Tdata_pthread (pid, tid, ptid,*pdata, * times)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_t *pdata;
pm_accu_time_t *times;
```

### Description

The **pm\_get\_data\_pthread** subroutine retrieves the current Performance Monitor data for a target pthread. The pthread must be stopped and must be part of a debuggee process under the control of a calling process.

The **pm\_get\_tdata\_pthread** subroutine retrieves the current Performance Monitor data for a target pthread, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata\_pthread** subroutine retrieves the current Performance Monitor data for a target pthread, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The Performance Monitor data is always a set (one per hardware counter on the machine used) of 64-bit values.

### Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
------------	--

<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data for the target pthread.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

<b>0</b>	No errors occurred.
<b>Positive error code</b>	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<b>/usr/include/pmapi.h</b>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_delete\_program\_pthread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pthread, pm\_get\_tdata\_pthread or pm\_get\_Tdata\_pthread Subroutine” on page 1054, “pm\_get\_program\_pthread Subroutine” on page 1077, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pthread Subroutine” on page 1093, “pm\_set\_program\_pthread Subroutine” on page 1113, “pm\_start\_pthread and pm\_tstart\_pthread Subroutine” on page 1127, “pm\_stop\_pthread and pm\_tstop\_pthread Subroutine ” on page 1136.

The **read\_real\_time** or **time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_pthread\_mx or pm\_get\_tdata\_pthread\_mx Subroutine

### Purpose

Returns Performance Monitor data in counter multiplexing mode for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>

int pm_get_data_thread_mx (pid, tid, ptid, *pmdata)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_mx_t *pmdata;

int pm_get_tdata_thread_mx (pid, tid, ptid, *pmdata, *time)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_data_mx_t *pmdata;
timebasestruct_t *time;
```

## Description

The **pm\_get\_data\_thread\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for a target pthread. The pthread must be stopped and must be part of a debuggee process under the control of a calling process.

The **pm\_get\_tdata\_thread\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for a target pthread, and a timestamp indicating the last time the hardware counters were read.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The Performance Monitor data is always an array of a set (one per hardware counter on the machine used) of 64-bit values.

The user application must free the array allocated to store accumulated counts and times ( the *accu\_set* field of the *pmdata* parameter).

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted) for the target pthread.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.



## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_thread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_thread\_mx and pm\_get\_program\_thread\_mm Subroutines” on page 1078, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_thread Subroutine” on page 1093, “pm\_set\_program\_thread\_mx and pm\_set\_program\_thread\_mm Subroutines” on page 1115, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1127, “pm\_stop\_thread and pm\_tstop\_thread Subroutine ” on page 1136.

The **read\_real\_time\_or\_time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine

### Purpose

Returns Performance Monitor data for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_data_thread (pid, tid, *pmdata)
pid_t pid;
tid_t tid;
pm_data_t *pmdata;
```

```
int pm_get_tdata_thread (pid, tid, *pmdata, *time)
pid_t pid;
tid_t tid;
pm_data_t *pmdata;
timebasestruct_t *time;
```

```
int pm_get_Tdata_thread (pid, tid, *pmdata, * times)
pm_data_t *pmdata;
pid_t pid;
tid_t tid;
pm_data_t *pmdata;
pm_accu_time_t *times;
```

## Description

These subroutines support only the 1:1 threading model. They have been superseded by the **pm\_get\_data\_pthread** and **pm\_get\_tdata\_pthread** subroutines, which support both the 1:1 and the M:N threading models. Calls to these subroutines are equivalent to calls to the **pm\_get\_data\_pthread** and **pm\_get\_tdata\_pthread** subroutines with a *ptid* parameter equal to 0.

The **pm\_get\_data\_thread** subroutine retrieves the current Performance Monitor data for a target kernel thread. The thread must be stopped and must be part of a debuggee process under the control of a calling process.

The **pm\_get\_tdata\_thread** subroutine retrieves the current Performance Monitor data for a target thread, and a timestamp indicating the last time the hardware counters were read.

The **pm\_get\_Tdata\_thread** subroutine retrieves the current Performance Monitor data for a target thread, and the accumulated time (timebase, PURR time and SPURR time) the events were counted.

The Performance Monitor data is always a set (one per hardware counter on the machine used) of 64-bit values.

## Parameters

<i>pid</i>	Process identifier of a target thread. The target process must be a debuggee of the caller process.
<i>tid</i>	Thread identifier of a target thread.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data for the target kernel thread.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.
<i>*times</i>	Pointer to a structure containing the accumulated time (timebase, PURR time and SPURR time) the events were counted. Each time counter can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_thread Subroutine” on page 1117, “pm\_get\_program\_thread Subroutine” on page 1080, “pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1057, “pm\_start\_thread and

pm\_tstart\_thread Subroutine” on page 1129, “pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137, “pm\_reset\_data\_thread Subroutine” on page 1094.

The **read\_real\_time** or **time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_data\_thread\_mx or pm\_get\_tdata\_thread\_mx Subroutine

### Purpose

Returns Performance Monitor data in counter multiplexing mode for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_get_data_thread_mx (pid, tid, *pdata)
pid_t pid;
tid_t tid;
pm_data_mx_t *pdata;

int pm_get_tdata_thread_mx (pid, tid, *pdata, *time)
pid_t pid;
tid_t tid;
pm_data_mx_t *pdata;
timebasestruct_t *time;
```

### Description

These subroutines support only the 1:1 threading model. They have been superseded by the **pm\_get\_data\_pthread\_mx** and **pm\_get\_tdata\_pthread\_mx** subroutines, which support both the 1:1 and the M:N threading models. Calls to these subroutines are equivalent to calls to the **pm\_get\_data\_pthread\_mx** and **pm\_get\_tdata\_pthread\_mx** subroutines with a *ptid* parameter equal to 0.

The **pm\_get\_data\_thread\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for a target kernel thread. The thread must be stopped and must be part of a debuggee process under the control of a calling process.

The **pm\_get\_tdata\_thread\_mx** subroutine retrieves the current Performance Monitor data in counter multiplexing mode for a target thread, and a timestamp indicating the last time the hardware counters were read.

The Performance Monitor data is always an array of a set (one per hardware counter on the machine used) of 64-bit values.

The user application must free the array allocated to store accumulated counts and times ( the *accu\_set* field of the *pdata* parameter).

### Parameters

<i>pid</i>	Process identifier of a target thread. The target process must be a debuggee of the caller process.
------------	---

<i>tid</i>	Thread identifier of a target thread.
<i>*pmdata</i>	Pointer to a structure to return the Performance Monitor data (array of accumulated counters, accumulated time and accumulated PURR and SPURR time for each event set counted) for the target kernel thread.
<i>*time</i>	Pointer to a structure containing the timebase value the last time the hardware Performance Monitoring counters were read. This can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	No errors occurred.
<b>Positive error code</b>	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_thread\_mx and pm\_set\_program\_thread\_mm Subroutines” on page 1119, “pm\_get\_program\_thread\_mx and pm\_get\_program\_thread\_mm Subroutines” on page 1082, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1129, “pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137, “pm\_reset\_data\_thread Subroutine” on page 1094.

The **read\_real\_time** or **time\_base\_to\_time** subroutine in in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_proctype Subroutine

### Purpose

Returns the current process type.

### Library

Performance Monitor APIs (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_proctype ()
```

## Description

The **pm\_get\_proctype** subroutine returns the current processor type. This value is the same as the one returned in the *proctype* parameter by the **pm\_initialize** subroutine.

## Return Values

Positive value	Current processor type.
-1	Unsupported processor type.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_initialize Subroutine” on page 1085.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program Subroutine

### Purpose

Retrieves systemwide Performance Monitor settings.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program ( *prog)  
pm_prog_t *prog;
```

### Description

The **pm\_get\_program** subroutine retrieves the current systemwide Performance Monitor settings. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers come from the lists returned by the **pm\_init** subroutine.

The counting mode includes user mode, the kernel mode, the current counting state, and the process tree mode. If the process tree mode is on, the counting applies only to the calling process and its descendants.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

If the events are represented by a group ID, then the **is\_group** bit is set in the mode, and the first element of the events array contains the group ID. The other elements of the events array are not meaningful.

## Parameters

*prog*

Returns which Performance Monitor events and modes are set. Supported modes are:

### PM\_USER

Counting processes running in user mode

### PM\_KERNEL

Counting processes running in kernel mode

### PM\_COUNT

Counting is on

### PM\_PROCTREE

Counting applies only to the calling process and its descendants

## Return Values

0

No errors occurred.

Positive error code

Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

*/usr/include/pmapi.h*

Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program Subroutine” on page 1095, “pm\_delete\_program Subroutine” on page 1028, “pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine” on page 1036, “pm\_start and pm\_tstart Subroutine” on page 1121, “pm\_stop and pm\_tstop Subroutine ” on page 1130, “pm\_reset\_data Subroutine” on page 1087.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_group Subroutine

### Purpose

Retrieves the Performance Monitor settings for the counting group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_group ( pid, tid, *prog)
```

```
pid_t pid;
tid_t tid;
pm_prog_t *prog;
```

## Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_get\_program\_pgroup** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_get\_program\_pgroup** subroutine with a *ptid* parameter equal to 0.

The **pm\_get\_program\_group** subroutine retrieves the Performance Monitor settings for the counting group to which a target kernel thread belongs. The thread must be stopped and must be part of a debuggee process under the control of the calling process. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers come from the lists returned by the **pm\_init** subroutine.

The counting mode includes the user mode and kernel mode, and the current counting state.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

## Parameters

<i>pid</i>	Process identifier of target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread identifier of the target thread.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. Supported modes are: <b>PM_USER</b> Counting process running in user mode <b>PM_KERNEL</b> Counting process running kernel mode <b>PM_COUNT</b> Counting is on <b>PM_PROCESS</b> Process level counting group

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_group Subroutine” on page 1096, “pm\_delete\_program\_group Subroutine” on page 1029, “pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine” on page 1039, “pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122, “pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131, “pm\_reset\_data\_group Subroutine” on page 1088.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_group\_mx and pm\_get\_program\_group\_mm Subroutines

### Purpose

Retrieves the Performance Monitor settings in counter multiplexing mode and multi-mode for the counting group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_group_mx ( pid, tid, *prog)
pid_t pid;
tid_t tid;
pm_prog_mx_t *prog;
```

```
int pm_get_program_group_mm ( pid, tid, *prog_mm)
pid_t pid;
tid_t tid;
pm_prog_mm_t *prog_mm;
```

### Description

These subroutines support only the 1:1 threading model. They have been superseded respectively by the **pm\_get\_program\_pgroup\_mx** subroutine and the **pm\_get\_program\_pgroup\_mm** subroutine, which support both the 1:1 and the M:N threading models. A call to the **pm\_get\_program\_group\_mx** subroutine or the **pm\_get\_program\_group\_mm** subroutine is respectively equivalent to a call to the **pm\_get\_program\_pgroup\_mx** subroutine or the **pm\_get\_program\_pgroup\_mm** subroutine with a *ptid* parameter equal to 0.

The **pm\_get\_program\_group\_mx** subroutine and the **pm\_get\_program\_group\_mm** subroutine retrieve the Performance Monitor settings for the counting group to which a target kernel thread belongs. The thread must be stopped and must be part of a debuggee process under the control of the calling process. This includes mode information and the events being counted, which are in an array of lists of event identifiers. The identifiers come from the lists returned by the **pm\_initialize** subroutine.

When counting in multiplexing mode (**pm\_get\_program\_group\_mx**), the mode is global to all of the events lists. When counting in multi-mode (**pm\_get\_program\_group\_mm**), a mode is associated with each event list.

Counting mode includes the user mode, the kernel mode, and the current counting state.



If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

The user application must free the allocated array to store the event lists (the *events\_set* field in the *prog* parameter).

## Parameters

<i>pid</i>	Process identifier of the target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread identifier of the target thread.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. It supports the following modes: <b>PM_USER</b> Counting process running in User Mode. <b>PM_KERNEL</b> Counting process running in Kernel Mode. <b>PM_COUNT</b> Counting is On. <b>PM_PROCESS</b> Process level counting group.
<i>*prog_mm</i>	Returns which Performance Monitor events and associated modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is on. <b>PM_PROCTREE</b> Counting that applies only to the calling process and its descendants.  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode are common to all modes set.

## Return Values

0	No errors occurred.
Positive error code	See the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

See the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_group\_mx and pm\_set\_program\_group\_mm Subroutines” on page 1098, “pm\_delete\_program\_group Subroutine” on page 1029, “pm\_get\_data\_group\_mx and

pm\_get\_tdata\_group\_mx Subroutine” on page 1041, “pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122, “pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131, and “pm\_reset\_data\_group Subroutine” on page 1088.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_mx and pm\_get\_program\_mm Subroutines

### Purpose

Retrieves system wide Performance Monitor settings in counter multiplexing mode and in multi-mode.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_mx ( *prog)  
pm_prog_mx_t *prog;
```

```
int pm_get_program_mm (*prog_mm)  
pm_prog_mm_t *prog_mm;
```

### Description

The **pm\_get\_program\_mx** and **pm\_get\_program\_mm** subroutines retrieve the current system wide Performance Monitor settings. This includes mode information and the events being counted, which are in an array of list of event identifiers. The identifiers come from the lists returned by the **pm\_initialize** subroutine. When you use the **pm\_get\_program\_mm** subroutine for multi-mode counting, a mode is associated to each event list.

The counting mode includes the user mode, the kernel mode, the current counting state, and the process tree mode. If the process tree mode is set, the counting applies only to the calling process and its descendants.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

If the events are represented by a group ID, then the **is\_group** bit is set in the mode, and the first element of each events array contains the group ID. The other elements of the events array are not used.

The user application must free the array allocated to store the event lists (events\_set field in prog).

## Parameters

*prog*

Returns which Performance Monitor events and modes are set. It supports the following modes:

**PM\_USER**

Counting processes running in the user mode.

**PM\_KERNEL**

Counting processes running in the kernel mode.

**PM\_COUNT**

Counting is on.

**PM\_PROCTREE**

Counting applies only to the calling process and its descendants.

*prog\_mm*

Returns which Performance Monitor events and associated modes are set. It supports the following modes:

**PM\_USER**

Counting processes running in the user mode.

**PM\_KERNEL**

Counting processes running in the kernel mode.

**PM\_COUNT**

Counting is On.

**PM\_PROCTREE**

Counting applies only to the calling process and its descendants.

The *PM\_PROCTREE* mode and the *PM\_COUNT* mode are common to all mode set.

## Return Values

0

No errors occurred.

**Positive error code**

Refer to the *pm\_error* (“*pm\_error* Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the “*pm\_error* Subroutine” on page 1035.

## Files

*/usr/include/pmapi.h*

Defines standard macros, data types, and subroutines.

## Related Information

The “*pm\_init* Subroutine” on page 1084, “*pm\_error* Subroutine” on page 1035, “*pm\_set\_program\_mx* and *pm\_set\_program\_mm* Subroutines” on page 1100, “*pm\_delete\_program* Subroutine” on page 1028, “*pm\_get\_data\_mx*, *pm\_get\_tdata\_mx*, *pm\_get\_data\_cpu\_mx*, *pm\_get\_tdata\_cpu\_mx*, *pm\_get\_data\_lcpu\_mx* and *pm\_get\_tdata\_lcpu\_mx* Subroutine” on page 1042, “*pm\_start* and *pm\_tstart* Subroutine” on page 1121, “*pm\_stop* and *pm\_tstop* Subroutine ” on page 1130, “*pm\_reset\_data* Subroutine” on page 1087.

## pm\_get\_program\_mygroup Subroutine

### Purpose

Retrieves the Performance Monitor settings for the counting group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_mygroup ( *prog)  
pm_prog_t *prog;
```

### Description

The **pm\_get\_program\_mygroup** subroutine retrieves the Performance Monitor settings for the counting group to which the calling kernel thread belongs. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers come from the lists returned by the **pm\_init** subroutine.

The counting mode includes user mode and kernel mode, and the current counting state.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

### Parameters

*\*prog*

Returns which Performance Monitor events and modes are set. Supported modes are:

**PM\_USER**

Counting processes running in user mode

**PM\_KERNEL**

Counting processes running in kernel mode

**PM\_COUNT**

Counting is on

**PM\_PROCESS**

Process level counting group

### Return Values

0

No errors occurred.

**Positive error code**

Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

### Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h`

Defines standard macros, data types, and subroutines.

## Related Information

The “`pm_init` Subroutine” on page 1084, “`pm_error` Subroutine” on page 1035, “`pm_set_program_mygroup` Subroutine” on page 1102, “`pm_delete_program_mygroup` Subroutine” on page 1030, “`pm_get_data_mygroup`, `pm_get_tdata_mygroup` or `pm_get_Tdata_mygroup` Subroutine” on page 1044, “`pm_start_mygroup` and `pm_tstart_mygroup` Subroutine” on page 1123, “`pm_stop_mygroup` and `pm_tstop_mygroup` Subroutine ” on page 1132, “`pm_reset_data_mygroup` Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## `pm_get_program_mygroup_mx` and `pm_get_program_mygroup_mm` Subroutines

### Purpose

Retrieves the Performance Monitor settings in counter multiplexing mode and multi-mode for the counting group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (`libpmapi.a`)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_mygroup_mx ( *prog)  
pm_prog_mx_t *prog;
```

```
int pm_get_program_mygroup_mm (*prog_mm)  
pm_prog_mm_t *prog_mm;
```

### Description

The `pm_get_program_mygroup_mx` and the `pm_get_program_mygroup_mm` subroutines retrieve the Performance Monitor settings for the counting group to which the calling kernel thread belongs. This includes mode information and the events being counted, which are in an array of lists of event identifiers. The identifiers come from the lists returned by the `pm_initialize` subroutine.

When counting in multiplexing mode, the mode is global to all of the events lists. When counting in multi-mode, a mode is associated to each event list.

Counting mode includes the user mode, the kernel mode, and the current counting state.

If the list includes an event that can be used with a threshold (as indicated by the `pm_init` subroutine), a threshold value is also returned.

The user application must free the allocated array to store the event lists ( the `events_set` field in the `prog` parameter).

## Parameters

<i>*prog</i>	Returns which Performance Monitor events and modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is on. <b>PM_PROCESS</b> Process level counting group.
<i>*prog_mm</i>	Returns which Performance Monitor events and associated modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is On. <b>PM_PROCTREE</b> Counting applies only to the calling processes and its descendants.  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode are common to all modes set.

## Return Values

0	No errors occurred.
Positive error code	See the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

See the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mygroup\_mx and pm\_set\_program\_mygroup\_mm Subroutines” on page 1103, “pm\_delete\_program\_mygroup Subroutine” on page 1030, “pm\_get\_data\_mygroup\_mx or pm\_get\_tdata\_mygroup\_mx Subroutine” on page 1046, “pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123, “pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132, and “pm\_reset\_data\_mygroup Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_mythread Subroutine

### Purpose

Retrieves the Performance Monitor settings for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_mythread ( *prog)  
pm_prog_t *prog;
```

### Description

The **pm\_get\_program\_mythread** subroutine retrieves the Performance Monitor settings for the calling kernel thread. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers come from the lists returned by the **pm\_init** subroutine.

The counting mode includes user mode and kernel mode, and the current counting state.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

### Parameters

*\*prog*

Returns which Performance Monitor events and modes are set. Supported modes are:

**PM\_USER**

Counting processes running in user mode

**PM\_KERNEL**

Counting processes running in kernel mode

**PM\_COUNT**

Counting is on

### Return Values

0

No errors occurred.

**Positive error code**

Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

### Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

### Files

**/usr/include/pmapi.h**

Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mythread Subroutine” on page 1106, “pm\_delete\_program\_mythread Subroutine” on page 1031, “pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine” on page 1047, “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133, “pm\_reset\_data\_mythread Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_mythread\_mx and pm\_get\_program\_mythread\_mm Subroutines

### Purpose

Retrieves the Performance Monitor settings in counter multiplexing mode and multi-mode for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_mythread_mx (*prog)  
pm_prog_mx_t *prog;
```

```
int pm_get_program_mythread_mm (*prog_mm)  
pm_prog_mm_t *prog_mm;
```

### Description

The **pm\_get\_program\_mythread\_mx** and the **pm\_get\_program\_mythread\_mm** subroutines retrieve the Performance Monitor settings for the calling kernel thread. This includes mode information and the events being counted, which are in an array of lists of event identifiers. The event identifiers come from the lists returned by the **pm\_initialize** subroutine.

When counting in multiplexing mode, the mode is global to all of the events lists. When counting in multi-mode, a mode is associated with each event list.

Counting mode includes the user mode, the kernel mode, and the current counting state.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

The user application must free the allocated array to store the event lists ( the `events_set` field in the `prog` parameter).



## Parameters

<i>*prog</i>	Returns which Performance Monitor events and modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is On.
<i>*prog_mm</i>	Returns which Performance Monitor events and associated modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is On. <b>PM_PROCTREE</b> Counting that applies only to the calling processes and its descendants.  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode are common to all modes set.

## Return Values

0	No errors occurred.
<b>Positive error code</b>	See the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

See the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mythread\_mx and pm\_set\_program\_mythread\_mm Subroutines” on page 1107, “pm\_delete\_program\_mythread Subroutine” on page 1031, “pm\_get\_data\_mythread\_mx or pm\_get\_tdata\_mythread\_mx Subroutine” on page 1049, “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133, and “pm\_reset\_data\_mythread Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_pgroup Subroutine

### Purpose

Retrieves Performance Monitor settings for the counting group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_pgroup ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_t *prog;
```

### Description

The **pm\_get\_program\_pgroup** subroutine retrieves the Performance Monitor settings for the counting group to which a target pthread belongs. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers come from the lists returned by the **pm\_initialize** subroutine.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The counting mode includes the user mode and kernel mode, and the current counting state.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value is also returned.

### Parameters

<i>pid</i>	Process ID of target pthread. The target process must be an argument of a debug process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. The following modes are supported: <b>PM_USER</b> Counts process running in user mode <b>PM_KERNEL</b> Counts process running kernel mode <b>PM_COUNT</b> Counting is on <b>PM_PROCESS</b> Process-level counting group

## Return Values

**0** No errors occurred.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_set\_program\_pgroup Subroutine” on page 1109, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_pgroup\_mx and pm\_get\_program\_pgroup\_mm Subroutines

### Purpose

Retrieves Performance Monitor settings in counter multiplexing mode and multi-mode for the counting group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_pgroup_mx ( pid, tid, ptid, *prog)  
pid_t pid;  
tid_t tid;  
ptid_t ptid;  
pm_prog_mx_t *prog;
```

```
int pm_get_program_pgroup_mm ( pid, tid, ptid, prog_mm)  
pid_t pid;  
tid_t tid;  
ptid_t ptid;  
pm_prog_mm_t *prog_mm;
```

### Description

The **pm\_get\_program\_pgroup\_mx** and the **pm\_get\_program\_pgroup\_mm** subroutine retrieve the Performance Monitor settings for the counting group to which a target pthread belongs. The pthread must

be stopped and must be part of a debuggee process, which is under the control of the calling process. This includes mode information and the events being counted, which are in an array of lists of event identifiers. The event identifiers come from the lists returned by the **pm\_initialize** subroutine.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with the *tid* parameter specified.

When counting in multiplexing mode, the mode is global to all of the events lists. When counting in the multi-mode, a mode is associated with each event list.

The counting mode includes the user mode and kernel mode, and the current counting state.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value is also returned.

The user application must free the allocated array to store the event lists (the *events\_set* field in the *prog* parameter).

## Parameters

<i>pid</i>	Process ID of target pthread. The target process must be an argument of a debug process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. It supports the following modes:  <b>PM_USER</b> Counts process running in User Mode.  <b>PM_KERNEL</b> Counts process running Kernel Mode.  <b>PM_COUNT</b> Counting is On.  <b>PM_PROCESS</b> Process-level counting group.
<i>*prog_mm</i>	Returns which Performance Monitor events and associated modes are set. It supports the following modes:  <b>PM_USER</b> Counting processes running in User Mode.  <b>PM_KERNEL</b> Counting processes running in Kernel Mode.  <b>PM_COUNT</b> Counting is On.  <b>PM_PROCTREE</b> Counting applies only to the calling processes and its descendants.  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode are common to all modes set.

## Return Values

<b>0</b>	No errors occurred.
<b>Positive error code</b>	See the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

See the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup\_mx and pm\_get\_tdata\_pgroup\_mx Subroutine” on page 1052, “pm\_set\_program\_pgroup\_mx and pm\_set\_program\_pgroup\_mm Subroutines” on page 1111, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, and “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_pthread Subroutine

### Purpose

Retrieves the Performance Monitor settings for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_pthread ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_t *prog;
```

### Description

The **pm\_get\_program\_pthread** subroutine retrieves the Performance Monitor settings for a target pthread. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The counting mode includes user mode and kernel mode, and the current counting state.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value is also returned.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be an argument of a debug process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. The following modes are supported: <b>PM_USER</b> Counts processes running in User Mode <b>PM_KERNEL</b> Counts processes running in Kernel Mode <b>PM_COUNT</b> Counting is On

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_delete\_program\_pthread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pthread, pm\_get\_tdata\_pthread or pm\_get\_Tdata\_pthread Subroutine” on page 1054, “pm\_set\_program\_pthread Subroutine” on page 1113, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pthread Subroutine” on page 1093, “pm\_start\_pthread and pm\_tstart\_pthread Subroutine” on page 1127, “pm\_stop\_pthread and pm\_tstop\_pthread Subroutine ” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_pthread\_mx and pm\_get\_program\_pthread\_mm Subroutines

### Purpose

Retrieves the Performance Monitor settings in counter multiplexing mode and multi-mode for a target pthread.

## Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>

int pm_get_program_thread_mx ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_mx_t *prog;

int pm_get_program_thread_mm ( pid, tid, ptid, prog_mm)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_mm_t *prog_mm;
```

## Description

The **pm\_get\_program\_thread\_mx** and the **pm\_set\_program\_thread\_mm** subroutines retrieve the Performance Monitor settings for a target pthread. The pthread must be stopped and must be part of a debuggee process, that is under the control of the calling process. This includes mode information and the events being counted, which are in an array of lists of event identifiers. The event identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

When counting in multiplexing mode, the mode is global to all of the events lists. When counting in the multi-mode, a mode is associated with each event list.

Counting mode includes the user mode, the kernel mode, and the current counting state.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value is also returned.

The user application must free the allocated array to store the event lists (the *events\_set* field in the *prog* parameter).

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be an argument of a debug process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. It supports the following modes: <b>PM_USER</b> Counts processes running in User Mode. <b>PM_KERNEL</b> Counts processes running in Kernel Mode. <b>PM_COUNT</b> Counting is On.

*\*prog\_mm* Returns which Performance Monitor events and associated modes are set. It supports the following modes:

**PM\_USER**

Counting processes running in User Mode.

**PM\_KERNEL**

Counting processes running in Kernel Mode.

**PM\_COUNT**

Counting is On.

**PM\_PROCTREE**

Counting that applies only to the calling processes and its descendants.

The *PM\_PROCTREE* mode and the *PM\_COUNT* mode are common to all modes set.

## Return Values

0 No errors occurred.

Positive error code See the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

See the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

*/usr/include/pmapi.h* Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_thread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_thread\_mx or pm\_get\_tdata\_thread\_mx Subroutine” on page 1055, “pm\_set\_program\_thread\_mx and pm\_set\_program\_thread\_mm Subroutines” on page 1115, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_thread Subroutine” on page 1093, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1127, and “pm\_stop\_thread and pm\_tstop\_thread Subroutine ” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_thread Subroutine

### Purpose

Retrieves the Performance Monitor settings for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_thread ( pid, tid, *prog)
```



```
pid_t pid;
tid_t tid;
pm_prog_t *prog;
```

## Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_get\_program\_pthread** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_get\_program\_pthread** subroutine with a *ptid* parameter equal to 0.

The **pm\_get\_program\_thread** subroutine retrieves the Performance Monitor settings for a target kernel thread. The thread must be stopped and must be part of a debuggee process under the control of the calling process. This includes mode information and the events being counted, which are in a list of event identifiers. The identifiers come from the lists returned by the **pm\_init** subroutine.

The counting mode includes user mode and kernel mode, and the current counting state.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

## Parameters

<i>pid</i>	Process identifier of the target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread identifier of the target thread.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. Supported modes are: <b>PM_USER</b> Counting processes running in <b>User</b> mode <b>PM_KERNEL</b> Counting processes running in <b>Kernel</b> mode <b>PM_COUNT</b> Counting is On

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_thread Subroutine” on page 1117, “pm\_delete\_program\_thread Subroutine” on page 1034, “pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1057, “pm\_start\_thread and

pm\_tstart\_thread Subroutine” on page 1129, “pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137, “pm\_reset\_data\_thread Subroutine” on page 1094.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_get\_program\_thread\_mx and pm\_get\_program\_thread\_mm Subroutines

### Purpose

Retrieves the Performance Monitor settings in counter multiplexing mode and multi-mode for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_get_program_thread_mx ( pid, tid, *prog)
pid_t pid;
tid_t tid;
pm_prog_mx_t *prog;
```

```
int pm_get_program_thread_mm (pid, tid, *prog_mm)
pid_t pid;
tid_t tid;
pm_prog_mm_t *prog_mm;
```

### Description

These subroutines support only the 1:1 threading model. They have been superseded respectively by the **pm\_get\_program\_pthread\_mx** and the **pm\_get\_program\_pthread\_mm** subroutines, which support both the 1:1 and the M:N threading models. A call to the **pm\_get\_program\_thread\_mx** subroutine or to the **pm\_get\_program\_thread\_mm** subroutine is respectively equivalent to a call to the **pm\_get\_program\_pthread\_mx** subroutine or the **pm\_get\_program\_pthread\_mm** subroutine with a *ptid* parameter equal to 0.

The **pm\_get\_program\_thread\_mx** subroutine and the **pm\_get\_program\_thread\_mm** subroutine retrieve the Performance Monitor settings for a target kernel thread. The thread must be stopped and must be part of a debuggee process under the control of the calling process. This includes mode information and the events being counted, which are in an array of list of event identifiers. The event identifiers come from the lists returned by the **pm\_initialize** subroutine.

When counting in multiplexing mode, the mode is global to all of the events lists. When counting in multi-mode, a mode is associated to each event list.

Counting mode includes the user mode, the kernel mode, and the current counting state.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value is also returned.

The user application must free the allocated array to store the event lists (the *events\_set* field in the *prog* parameter).

## Parameters

<i>pid</i>	Process identifier of the target thread. The target process must be an argument of a debug process.
<i>tid</i>	Thread identifier of the target thread.
<i>*prog</i>	Returns which Performance Monitor events and modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is On.
<i>*prog_mm</i>	Returns which Performance Monitor events and associated modes are set. It supports the following modes: <b>PM_USER</b> Counting processes running in User Mode. <b>PM_KERNEL</b> Counting processes running in Kernel Mode. <b>PM_COUNT</b> Counting is On. <b>PM_PROCTREE</b> Counting that applies only to the calling process and its descendants.  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode are common to all modes set.

## Return Values

0	No errors occurred.
Positive error code	See the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

See the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_thread\_mx and pm\_set\_program\_thread\_mm Subroutines” on page 1119, “pm\_delete\_program\_thread Subroutine” on page 1034, “pm\_get\_data\_thread\_mx or pm\_get\_tdata\_thread\_mx Subroutine” on page 1059, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1129, “pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137, and “pm\_reset\_data\_thread Subroutine” on page 1094.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_init Subroutine

### Purpose

Initializes the Performance Monitor APIs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_init ( filter, *pminfo, *pm_groups_info)
int filter;
pm_info_t *pminfo;
pm_groups_info_t *pm_groups_info;
```

### Description

**Note:** The **pm\_init** subroutine cannot be used on processors newer than POWER4. With such processors, the **pm\_initialize** subroutine must be used.

The **pm\_init** subroutine initializes the Performance Monitor API library. It returns, after taking into account a *filter* on the status of the events, the number of counters available on this processor, and one table per counter with the list of events available. For each event, an event identifier, a status, a flag indicating if the event can be used with a threshold, two names, and a description are provided.

The event identifier is used with all the **pm\_set\_program** interfaces and is also returned by all of the **pm\_get\_program** interfaces. Only event identifiers present in the table returned can be used. In other words, the *filter* is effective for all API calls.

The status describes whether the event has been verified, is still unverified, or works with some caveat, as explained in the description. This field is necessary because the filter can be any combination of the three available status bits. The flag points to events that can be used with a threshold.

Only events categorized as *verified* have gone through full verification. Events categorized as *caveat* have been verified only within the limitations documented in the event description. Events categorized as *unverified* have undefined accuracy. Use caution with *unverified* events; the Performance Monitor software is essentially providing a service to read hardware registers which may or may not have any meaningful content. Users may experiment with unverified event counters and determine for themselves what, if any, use they may have for specific tuning situations.

The short mnemonic name is provided for easy keyword-based search in the event table (see the sample program **/usr/samples/pmapi/sysapit2.c** for code using mnemonic names). The complete name of the event is also available and a full description for each event is returned.

The structure returned also has the threshold multiplier for this processor and the processor name

On some platforms, it is possible to specify event groups instead of individual events. Event groups are predefined sets of events. Rather than specify each event individually, a single group ID is specified. On some platforms, such as POWER4, use of the event groups is required, and attempts to specify individual events return an error.

The interface to **pm\_init** has been enhanced to return the list of supported event groups in an optional third parameter. For binary compatibility, the third parameter must be explicitly requested by OR-ing the bitflag, **PM\_GET\_GROUPS**, into the *filter* parameter.

If the *pm\_groups\_info* parameter returned by **pm\_init** is **NULL**, there are no supported event groups for the platform. Otherwise an array of **pm\_groups\_t** structures are returned in the **event\_groups** field. The length of the array is given by the **max\_groups** field.

The **pm\_groups\_t** structure contains a group identifier, two names and a description that are similar to those of the individual events. In addition, there is an array of integers that specify the events contained in the group.

## Parameters

<i>filter</i>	Specifies which event types to return. <b>PM_VERIFIED</b> Events which have been verified <b>PM_UNVERIFIED</b> Events which have not been verified <b>PM_CAVEAT</b> Events which are usable but with caveats as described in the long description
<i>*pminfo</i>	Returned structure with processor name, threshold multiplier, and a filtered list of events with their current status.
<i>*pm_groups_info</i>	Returned structure with list of supported groups. This parameter is only meaningful if <b>PM_GET_GROUPS</b> is OR-ed into the <i>filter</i> parameter.

## Return Values

<b>0</b>	No errors occurred.
<b>Positive error code</b>	Refer to the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

See the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

“pm\_initialize Subroutine.”

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_initialize Subroutine

### Purpose

Initializes the Performance Monitor APIs and returns information about a processor.

## Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>

int pm_initialize ( filter, *pminfo, *pmgroups, proctype)
int filter;
pm_info2_t *pminfo;
pm_groups_info_t *pmgroups;
int proctype;
```

## Description

The **pm\_initialize** subroutine initializes the Performance Monitor API library and retrieves information about a type of processor (if the specified *proctype* is not **PM\_CURRENT**). It takes into account a *filter* on the events status, then it returns the number of counters available on this processor and one table per counter containing the list of available events. For each event, it provides an event identifier, a status, two names, and a description. The status contains a set of flags indicating: the event status, if the event can be used with a threshold, if the event is a shared event, and if the event is a grouped-only event.

The event identifier is used with all **pm\_set\_program** interfaces and is also returned by all of the **pm\_get\_program** interfaces. Only event identifiers present in the returned table can be used. In other words, the *filter* is effective for all API calls.

The status describes whether the event has been verified, is still unverified, or works with some caveat, as explained in the description. This field is necessary because the filter can be any combination of the three available status bits. The flag points to events that can be used with a threshold.

Only events categorized as verified have been fully verified. Events categorized as *caveat* have been verified only with the limitations documented in the event description. Events categorized as *unverified* have an undefined accuracy. Use *unverified* events cautiously; the Performance Monitor software provides essentially a service to read hardware registers, which might or might not have meaningful content. Users might experiment for themselves with unverified event counters to determine if they can be used for specific tuning situations.

The short mnemonic name is provided for an easy keyword-based search in the event table (see the sample program **/usr/samples/pmapi/cpi.c** for code using mnemonic names). The complete name of the event is also available, and a full description for each event is returned.

The returned structure also contains the threshold multipliers for this processor, the processor name, and its characteristics. On some platforms, up to three threshold multipliers are available.

On some platforms, it is possible to specify event groups instead of individual events. Event groups are predefined sets of events. Rather than specify each event individually, a single group ID is specified. On some platforms, such as POWER4, using event groups is mandatory, and specifying individual events returns an error.

The interface to **pm\_initialize** returns the list of supported event groups in its third parameter. If the *pmgroups* parameter returned by **pm\_initialize** is NULL, there are no supported event groups for the platform. Otherwise an array of **pm\_groups\_t** structures is returned in the **event\_groups** field. The length of the array is given by the **max\_groups** field.

The **pm\_groups\_t** structure contains a group identifier, two names, and a description that are all similar to those of the individual events. In addition, an array of integers specifies the events contained in the group.

If the *proctype* parameter is not set to **PM\_CURRENT**, the Performance Monitor APIs library is not initialized, and the subroutine only returns information about the specified processor and those events and groups available in its parameters (*pminfo* and *pmgroups*) taking into account the filter. If the *proctype* parameter is set to **PM\_CURRENT**, in addition to returning the information described, the Performance Monitor APIs library is initialized and ready to accept other calls.

## Parameters

<i>filter</i>	Specifies which event types to return. <b>PM_VERIFIED</b> Events that have been verified. <b>PM_UNVERIFIED</b> Events that have not been verified. <b>PM_CAVEAT</b> Events that are usable but with caveats, as explained in the long description.
<i>pmgroups</i>	Returned structure containing the list of supported groups.
<i>pminfo</i>	Returned structure containing the processor name, the threshold multiplier and a filtered list of events with their current status.
<i>proctype</i>	Initializes the Performance Monitor API and retrieves information about a specific processor type: <b>PM_CURRENT</b> Retrieves information about the current processor and initializes the Performance Monitor API library. <b>other</b> Retrieves information about a specific processor.

## Return Values

0	No errors occurred.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The **pm\_initialize** subroutine replaces **pm\_init** subroutine. It is mandatory to initialize the Performance Monitor API library for processors newer than POWER4.

“pm\_error Subroutine” on page 1035.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data Subroutine

### Purpose

Resets system wide Performance Monitor data.

## Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>

int pm_reset_data ()
```

## Description

The **pm\_reset\_data** subroutine resets the current system wide Performance Monitor data. The data is a set (one per hardware counter on the machine used) of 64-bit values. All values are reset to 0.

## Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

See the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_set\_program** (“pm\_set\_program Subroutine” on page 1095) subroutine, **pm\_get\_program** (“pm\_get\_program Subroutine” on page 1061) subroutine, **pm\_delete\_program** (“pm\_delete\_program Subroutine” on page 1028) subroutine, **pm\_get\_data** (“pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine” on page 1036) subroutine, **pm\_start** (“pm\_start and pm\_tstart Subroutine” on page 1121) subroutine, **pm\_stop** (“pm\_stop and pm\_tstop Subroutine ” on page 1130) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data\_group Subroutine

### Purpose

Resets Performance Monitor data for a target thread and the counting group to which it belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```



```
int pm_reset_data_group ( pid, tid)
pid_t pid;
tid_t tid;
```

## Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_reset\_data\_pgroup** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_reset\_data\_pgroup** subroutine with a *ptid* parameter equal to 0.

The **pm\_reset\_data\_group** subroutine resets the current Performance Monitor data for a target kernel thread and the counting group to which it belongs. The thread must be stopped and must be part of a debuggee process, under control of the calling process. The data is a set (one per hardware counter on the machine used) of 64-bit values. All values are reset to 0. Because the data for all the other threads in the group is not affected, the group is marked as inconsistent unless it has only one member.

## Parameters

<i>pid</i>	Process ID of target thread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target thread.

## Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_set\_program\_group** (“pm\_set\_program\_group Subroutine” on page 1096) subroutine, **pm\_get\_program\_group** (“pm\_get\_program\_group Subroutine” on page 1062) subroutine, **pm\_delete\_program\_group** (“pm\_delete\_program\_group Subroutine” on page 1029) subroutine, **pm\_start\_group** (“pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122) subroutine, **pm\_stop\_group** (“pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131) subroutine, **pm\_get\_data\_group** (“pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine” on page 1039) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data\_mygroup Subroutine

### Purpose

Resets Performance Monitor data for the calling thread and the counting group to which it belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_reset_data_mygroup()
```

### Description

The **pm\_reset\_data\_mygroup** subroutine resets the current Performance Monitor data for the calling kernel thread and the counting group to which it belongs. The data is a set (one per hardware counter on the machine used) of 64-bit values. All values are reset to 0. Because the data for all the other threads in the group is not affected, the group is marked as inconsistent unless it has only one member.

### Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

### Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

### Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

### Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_set\_program\_mygroup** (“pm\_set\_program\_mygroup Subroutine” on page 1102) subroutine, **pm\_get\_program\_mygroup** (“pm\_get\_program\_mygroup Subroutine” on page 1068) subroutine, **pm\_delete\_program\_mygroup** (“pm\_delete\_program\_mygroup Subroutine” on page 1030) subroutine, **pm\_start\_mygroup** (“pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123) subroutine, **pm\_stop\_mygroup** (“pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132) subroutine, **pm\_get\_data\_mygroup** (“pm\_get\_data\_mygroup, pm\_get\_tdata\_mygroup or pm\_get\_Tdata\_mygroup Subroutine” on page 1044) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data\_mythread Subroutine

### Purpose

Resets Performance Monitor data for the calling thread.

## Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>
```

```
int pm_reset_data_mythread()
```

## Description

The **pm\_reset\_data\_mythread** subroutine resets the current Performance Monitor data for the calling kernel thread. The data is a set (one per hardware counter on the machine) of 64-bit values. All values are reset to 0.

## Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_set\_program\_mythread** (“pm\_set\_program\_mythread Subroutine” on page 1106) subroutine, **pm\_get\_program\_mythread** (“pm\_get\_program\_mythread Subroutine” on page 1071) subroutine, **pm\_delete\_program\_mythread** (“pm\_delete\_program\_mythread Subroutine” on page 1031) subroutine, **pm\_start\_mythread** (“pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125) subroutine, **pm\_stop\_mythread** (“pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133) subroutine, **pm\_get\_data\_mythread** (“pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine” on page 1047) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data\_pgroup Subroutine

### Purpose

Resets Performance Monitor data for a target pthread and the counting group to which it belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_reset_data_pgroup ( pid, tid, ptid)
pid_t pid;
tid_t tid;
ptid_t ptid;
```

## Description

The **pm\_reset\_data\_pgroup** subroutine resets the current Performance Monitor data for a target pthread and the counting group to which it belongs. The pthread must be stopped and must be part of a debuggee process, under control of the calling process. The data is a set (one per hardware counter on the machine used) of 64-bit values. All values are reset to 0. Because the data for all the other pthreads in the group is not affected, the group is marked as inconsistent unless it has only one member.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_get\_program\_pgroup Subroutine” on page 1074, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_set\_program\_pgroup Subroutine” on page 1109, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data\_pthread Subroutine

### Purpose

Resets Performance Monitor data for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_reset_data_pthread ( pid, tid, ptid)
pid_t pid;
tid_t tid;
ptid_t ptid;
```

### Description

The **pm\_reset\_data\_pthread** subroutine resets the current Performance Monitor data for a target pthread. The pthread must be stopped and must be part of a debuggee process. The data is a set (one per hardware counter on the machine used) of 64-bit values. All values are reset to 0.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

### Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.

### Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

### Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

### Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

### Related Information

The “pm\_delete\_program\_pthread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pthread, pm\_get\_tdata\_pthread or pm\_get\_Tdata\_pthread Subroutine” on page 1054,

“pm\_get\_program\_pthread Subroutine” on page 1077, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pthread Subroutine” on page 1093, “pm\_set\_program\_pthread Subroutine” on page 1113, “pm\_start\_pthread and pm\_tstart\_pthread Subroutine” on page 1127, “pm\_stop\_pthread and pm\_tstop\_pthread Subroutine” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_reset\_data\_thread Subroutine

### Purpose

Resets Performance Monitor data for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_reset_data_thread ( pid, tid)
pid_t pid;
tid_t tid;
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_reset\_data\_pthread** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_reset\_data\_pthread** subroutine with a *ptid* parameter equal to 0.

The **pm\_reset\_data\_thread** subroutine resets the current Performance Monitor data for a target kernel thread. The thread must be stopped and must be part of a debuggee process. The data is a set (one per hardware counter on the machine used) of 64-bit values. All values are reset to 0.

### Parameters

<i>pid</i>	Process id of target thread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread id of target thread.

### Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

### Error Codes

Refer to the **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine.

### Files

<b>/usr/include/pmapi.h</b>	Defines standard macros, datatypes, and subroutines.
-----------------------------	--

## Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_set\_program\_thread** (“pm\_set\_program\_thread Subroutine” on page 1117) subroutine, **pm\_get\_program\_thread** (“pm\_get\_program\_thread Subroutine” on page 1080) subroutine, **pm\_delete\_program\_thread** (“pm\_delete\_program\_thread Subroutine” on page 1034) subroutine, **pm\_start\_thread** (“pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1129) subroutine, **pm\_stop\_thread** (“pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137) subroutine, **pm\_get\_data\_thread** (“pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1057) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program Subroutine

### Purpose

Sets system wide Performance Monitor programming.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program ( *prog)  
pm_prog_t *prog;
```

### Description

The **pm\_set\_program** subroutine sets system wide Performance Monitor programming. The setting includes the events to be counted, and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_init** subroutine.

The counting mode includes User Mode and/or Kernel Mode, the Initial Counting State, and the Process Tree Mode. The Process Tree Mode sets counting to On only for the calling process and its descendants. The defaults are set to Off for User Mode and Kernel Mode. The initial default state is set to delay counting until the **pm\_start** subroutine is called, and to count the activity of all the processes running in the system.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

On some platforms, event groups can be specified instead of individual events. This is done by setting the bitfield **is\_group** in the mode, and placing the group ID into the first element of the events array. (The group ID was obtained by **pm\_init**).

## Parameters

<i>*prog</i>	Specifies the events and modes to use in Performance Monitor setup. The following modes are supported:
<i>PM_USER</i>	Counts processes running in User Mode (default is set to Off)
<i>PM_KERNEL</i>	Counts processes running in Kernel Mode (default is set to Off)
<i>PM_COUNT</i>	Starts counting immediately (default is set to Not to Start Counting)
<i>PM_PROCTREE</i>	Sets counting to On only for the calling process and its descendants (default is set to Off)

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program Subroutine” on page 1061, “pm\_delete\_program Subroutine” on page 1028, “pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine” on page 1036, “pm\_start and pm\_tstart Subroutine” on page 1121, “pm\_stop and pm\_tstop Subroutine ” on page 1130, “pm\_reset\_data Subroutine” on page 1087.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_group Subroutine

### Purpose

Sets Performance Monitor programming for a target thread and creates a counting group.

### Library

Performance Monitor APIs Library (**libpmapi.a**)



## Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_group ( pid, tid, *prog)
pid_t pid;
tid_t tid;
pm_prog_t *prog;
```

## Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_set\_program\_pgroup** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_set\_program\_pgroup** subroutine with a *ptid* parameter equal to 0.

The **pm\_set\_program\_group** subroutine sets the Performance Monitor programming for a target kernel thread. The thread must be stopped and must be part of a debuggee process, under the control of the calling process. The setting includes the events to be counted and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_init** subroutine.

This call also creates a counting group, which includes the target thread and any thread which it, or any of its descendants, will create in the future. Optionally, the group can be defined as also containing all the existing and future threads belonging to the target process.

The counting mode includes User Mode and/or Kernel Mode, and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_group** subroutine is called.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

## Parameters

*pid*

Process ID of target thread. Target process must be a debuggee of a calling process.

*tid*

Thread ID of target thread.

*\*prog*

### **PM\_USER**

Counts processes running in User Mode (default is set to Off)

### **PM\_KERNEL**

Counts processes running in Kernel Mode (default is set to Off)

### **PM\_COUNT**

Starts counting immediately (default is set to Not to Start Counting)

### **PM\_PROCESS**

Creates a process-level counting group

## Return Values

0

Operation completed successfully.

**Positive error code**

Refer to the "pm\_error Subroutine" on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_group Subroutine” on page 1062, “pm\_delete\_program\_group Subroutine” on page 1029, “pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine” on page 1039, “pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122, “pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131, “pm\_reset\_data\_group Subroutine” on page 1088.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_group\_mx and pm\_set\_program\_group\_mm Subroutines

### Purpose

Sets the Performance Monitor program in counter multiplexing mode and multi-mode for a target thread and creates a counting group.

### Library

Performance Monitor APIs Library (`libpmapi.a`)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_group_mx ( pid, tid, *prog)  
pid_t pid;  
tid_t tid;  
pm_prog_mx_t *prog;
```

```
int pm_set_program_group_mm ( pid, tid, *prog_mm)  
pid_t pid;  
tid_t tid;  
pm_prog_mm_t *prog_mm;
```

### Description

The `pm_set_program_group_mx` and `pm_set_program_group_mm` subroutines support only the 1:1 threading model. They have been superseded respectively by the `pm_set_program_pgroup_mx` and `pm_set_program_pgroup_mm` subroutines, which support both the 1:1 and the M:N threading models. A call to the `pm_set_program_pgroup_mx` or `pm_set_program_pgroup_mm` subroutine is respectively equivalent to a call to the `pm_set_program_group_mx` or `pm_set_program_group_mm` subroutine with a `ptid` parameter equal to 0.

The `pm_set_program_group_mx` and `pm_set_program_group_mm` subroutines set the Performance Monitor program respectively in counter multiplexing mode or in multi-mode for a target kernel thread. The thread must be stopped and must be part of a debuggee process, which is under the control of the calling process.

The **pm\_set\_program\_group\_mx** subroutine setting includes the list of the event arrays to be counted and the mode in which to count. The mode is global to all of the event lists. The events to count are in an array of lists of event identifiers.

The **pm\_set\_program\_group\_mm** subroutine setting includes the list of the event arrays to be counted, and the associated mode in which to count each event array. A counting mode is defined for each event array.

The event identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

Both subroutines create a counting group, which includes the target thread and any thread which it, or any of its descendants, will create in the future. The group can also be defined as containing all the existing and future threads belonging to the target process.

The counting mode for the subroutines includes the User Mode, the Kernel Mode, or both of them, and the Initial Counting State. The default is set to Off for the User Mode and the Kernel Mode. The initial default state is set to delay counting until the **pm\_start\_group** subroutine is called.

When you use the **pm\_set\_program\_group\_mm** subroutine for multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by their values that are defined in the first programming set.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Specifies the process ID of target thread. The target process must be a debuggee of a calling process.
<i>tid</i>	Specifies the thread ID of the target thread.
<i>*prog</i>	Specifies the events and modes to use in the Performance Monitor setup. The <i>prog</i> parameter supports the following modes: <b>PM_USER</b> Counts processes running in User Mode (default is set to Off). <b>PM_KERNEL</b> Counts processes running in Kernel Mode (default is set to Off). <b>PM_COUNT</b> Starts counting immediately (default is set to Not to start counting). <b>PM_PROCESS</b> Creates a process-level counting group.
<i>* prog_mm</i>	Specifies the events and the modes to use in the Performance Monitor setup. The <i>prog_mm</i> parameter supports the following modes: <b>PM_USER</b> Counts processes running in User Mode (default is set to Off). <b>PM_KERNEL</b> Counts processes running in Kernel Mode (default is set to Off). <b>PM_COUNT</b> Starts counting immediately (default is set to Not to start counting). <b>PM_PROCTREE</b> Sets counting to On only for the calling process and its descendants (default is set to Off).  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode defined in the first setting fix value for the counting.

## Return Values

0	Operation completed successfully.
Positive Error Code	See the <b>pm_error</b> (“pm_error Subroutine” on page 1035) subroutine to decode the error code.

## Error Codes

See the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_group\_mx and pm\_get\_program\_group\_mm Subroutines” on page 1064, “pm\_delete\_program\_group Subroutine” on page 1029, “pm\_get\_data\_group\_mx and pm\_get\_tdata\_group\_mx Subroutine” on page 1041, “pm\_start\_group and pm\_tstart\_group Subroutine” on page 1122, “pm\_stop\_group and pm\_tstop\_group Subroutine ” on page 1131, and “pm\_reset\_data\_group Subroutine” on page 1088.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_mx and pm\_set\_program\_mm Subroutines

### Purpose

Sets system wide Performance Monitor programming in counter multiplexing mode and in multi-mode.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_set_program_mx (*prog)
pm_prog_mx_t *prog;

int pm_set_program_mm (*prog_mm)
pm_prog_mm_t *prog_mm;
```

### Description

The **pm\_set\_program\_mx** and **pm\_set\_program\_mm** subroutines set system wide Performance Monitor programming in counter multiplexing mode.

The **pm\_set\_program\_mx** setting includes the list of the event arrays to be counted, and a mode in which to count. The events to count are in an array of list of event identifiers. The mode is global to all the event lists.

The **pm\_set\_program\_mm** setting includes the list of the event arrays to be counted, and the associated mode in which to count each event array. A counting mode is defined for each event array.

The identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

The counting mode includes the User Mode and the Kernel Mode, or either of them; the Initial Counting State; and the Process Tree Mode. The Process Tree Mode sets counting to On only for the calling process and its descendants. The defaults are set to Off for the User Mode and the Kernel Mode. The initial default state is set to delay counting until the **pm\_start** subroutine is called, and to count the activity of all the processes running in the system.

When you use the **pm\_set\_program\_mm** subroutine for multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by their values that are defined in the first programming set.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

On some platforms, event groups can be specified instead of individual events. This is done by setting the **is\_group** bitfield in the mode, and placing the group ID into the first element of each events array. (The group ID was obtained by **pm\_init** subroutine.)

## Parameters

*\*prog*

Specifies the events and modes to use in Performance Monitor setup. It supports the following modes:

*PM\_USER*

Counts processes that run in the User Mode (default is set to Off).

*PM\_KERNEL*

Counts processes that run in the Kernel Mode (default is set to Off).

*PM\_COUNT*

Starts counting immediately (default is set to Not to Start Counting).

*PM\_PROCTREE*

Sets counting to On only for the calling process and its descendants (default is set to Off).

*\*prog\_mm*

Specifies the events and the associated modes to use in the Performance Monitor setup. It supports the following modes:

*PM\_USER*

Counts processes that run in the User Mode (default is set to Off).

*PM\_KERNEL*

Counts processes that run in the Kernel Mode (default is set to Off).

*PM\_COUNT*

Starts counting immediately (default is set to Not to start counting).

*PM\_PROCTREE*

Sets counting to On only for the calling process and its descendants (default is set to Off).

The *PM\_PROCTREE* and the *PM\_COUNT* modes defined in the first setting fix the value for the counting.

## Return Values

0 Operation completed successfully.  
Positive Error Code Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_mx and pm\_get\_program\_mm Subroutines” on page 1066, “pm\_get\_program\_mx and pm\_get\_program\_mm Subroutines” on page 1066, “pm\_delete\_program Subroutine” on page 1028, “pm\_get\_data\_mx, pm\_get\_tdata\_mx, pm\_get\_data\_cpu\_mx, pm\_get\_tdata\_cpu\_mx, pm\_get\_data\_lcpu\_mx and pm\_get\_tdata\_lcpu\_mx Subroutine” on page 1042, “pm\_start and pm\_tstart Subroutine” on page 1121, “pm\_stop and pm\_tstop Subroutine ” on page 1130, “pm\_reset\_data Subroutine” on page 1087.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_mygroup Subroutine

### Purpose

Sets Performance Monitor programming for the calling thread and creates a counting group.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_mygroup ( *prog)  
pm_prog_t *prog;
```

### Description

The **pm\_set\_program\_mygroup** subroutine sets the Performance Monitor programming for the calling kernel thread. The setting includes the events to be counted and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_init** subroutine.

This call also creates a counting group, which includes the calling thread and any thread which it, or any of its descendants, will create in the future. Optionally, the group can be defined as also containing all the existing and future threads belonging to the calling process.

The counting mode includes User Mode and/or Kernel Mode, and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_mygroup** subroutine is called.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

## Parameters

*\*prog*

Specifies the events and mode to use in Performance Monitor setup. The following modes are supported:

### **PM\_USER**

Counts processes running in User Mode (default is set to Off)

### **PM\_KERNEL**

Counts processes running in Kernel Mode (default is set to Off)

### **PM\_COUNT**

Starts counting immediately (default is set to Not to Start Counting)

### **PM\_PROCESS**

Creates a process-level counting group

## Return Values

0

Operation completed successfully.

**Positive error code**

Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h**

Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_mygroup Subroutine” on page 1068, “pm\_delete\_program\_mygroup Subroutine” on page 1030, “pm\_get\_data\_mygroup, pm\_get\_tdata\_mygroup or pm\_get\_Tdata\_mygroup Subroutine” on page 1044, “pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123, “pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132, “pm\_reset\_data\_mygroup Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## **pm\_set\_program\_mygroup\_mx and pm\_set\_program\_mygroup\_mm Subroutines**

### **Purpose**

Sets Performance Monitor programming in counter multiplexing mode and multi-mode for the calling thread and creates a counting group.

## Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_mygroup_mx ( *prog)  
pm_prog_mx_t *prog;
```

```
int pm_set_program_mygroup_mm (*prog_mm)  
pm_prog_mm_t *prog_mm;
```

## Description

The **pm\_set\_program\_mygroup\_mx** and **pm\_set\_program\_mygroup\_mm** subroutines set the Performance Monitor programming respectively in counter multiplexing mode or in multi-mode for the calling kernel thread.

The **pm\_set\_program\_mygroup\_mx** subroutine setting includes the list of event arrays to be counted and a mode in which to count. The mode is global to all of the event lists. The events to count are in an array of list of event identifiers.

The **pm\_set\_program\_mygroup\_mm** subroutine setting includes the list of the event arrays to be counted, and the mode in which to count each event array. A counting mode is defined for each event array.

The identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

Both subroutines create a counting group, which includes the calling thread and any thread which it, or any of its descendants, will create in the future. Optionally, the group can be defined as also containing all the existing and future threads belonging to the calling process.

The counting mode for both subroutines includes the User Mode or the Kernel Mode, or both of them; the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_mygroup** subroutine is called.

When you use the **pm\_set\_program\_mygroup\_mm** subroutine for multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by their values defined in the first programming set.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.



## Parameters

*\*prog*

Specifies the events and modes to use in Performance Monitor setup. The *prog* parameter supports the following modes:

### **PM\_USER**

Counts processes running in User Mode (default is set to Off).

### **PM\_KERNEL**

Counts processes running in Kernel Mode (default is set to Off).

### **PM\_COUNT**

Starts counting immediately (default is set to Not to Start Counting).

### **PM\_PROCESS**

Creates a process-level counting group.

*\*prog\_mm*

Specifies the events and the associated modes to use in the Performance Monitor setup. The *prog\_mm* parameter supports the following modes:

### **PM\_USER**

Counts processes running in the User Mode (default is set to Off).

### **PM\_KERNEL**

Counts processes running in the Kernel Mode (default is set to Off).

### **PM\_COUNT**

Starts counting immediately (default is set to Not to start counting).

### **PM\_PROCTREE**

Sets counting to On only for the calling process and its descendants (default is set to Off).

The *PM\_PROCTREE* mode and the *PM\_COUNT* mode defined in the first setting fix the value for the counting.

## Return Values

0

Operation completed successfully.

Positive Error Code

Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

*/usr/include/pmapi.h*

Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_mygroup\_mx and pm\_get\_program\_mygroup\_mm Subroutines” on page 1069, “pm\_delete\_program\_mygroup Subroutine” on page 1030, “pm\_get\_data\_mygroup\_mx or

pm\_get\_tdata\_mygroup\_mx Subroutine” on page 1046, “pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine” on page 1123, “pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132, “pm\_reset\_data\_mygroup Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_mythread Subroutine

### Purpose

Sets Performance Monitor programming for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_mythread ( *prog)  
pm_prog_t *prog;
```

### Description

The **pm\_set\_program\_mythread** subroutine sets the Performance Monitor programming for the calling kernel thread. The setting includes the events to be counted, and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_init** subroutine.

The counting mode includes User Mode and/or Kernel Mode, and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_mythread** subroutine is called.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

### Parameters

*\*prog*

Specifies the event modes to use in Performance Monitor setup. The following modes are supported:

**PM\_USER**

Counts processes running in User Mode (default is set to Off)

**PM\_KERNEL**

Counts processes running in Kernel Mode (default is set to Off)

**PM\_COUNT**

Starts counting immediately (default is set to Not to Start Counting)

**PM\_PROCESS**

Creates a process-level counting group

## Return Values

0 Operation completed successfully.  
Positive error code Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_mythread Subroutine” on page 1071, “pm\_delete\_program\_mythread Subroutine” on page 1031, “pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine” on page 1047, “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133, “pm\_reset\_data\_mythread Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_mythread\_mx and pm\_set\_program\_mythread\_mm Subroutines

### Purpose

Sets Performance Monitor programming in counter multiplexing mode and multi-mode for the calling thread.

### Library

Performance Monitor APIs Library (`libpmapi.a`)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_mythread_mx ( *prog)  
pm_prog_mx_t *prog;
```

```
int pm_set_program_mythread_mm ( *prog_mm)  
pm_prog_mm_t *prog_mm;
```

### Description

The `pm_set_program_mythread_mx` and the `pm_set_program_mythread_mm` subroutines set the Performance Monitor programming respectively in counter multiplexing mode or in multi-mode for the calling kernel thread.

The `pm_set_program_mythread_mx` subroutine setting includes the list of the event arrays to be counted, and a mode in which to count. The mode is global to all event lists. The events to count are in an array of list of event identifiers.

The **pm\_set\_program\_mythread\_mm** setting includes the lists of the event arrays to be counted, and the associated modes in which to count each event array. A counting mode is defined for each event array.

The event identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

The counting mode for both subroutines includes the User Mode or the Kernel Mode, or both of them; and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_mythread** subroutine is called.

When you use the **pm\_set\_program\_mythread\_mm** subroutine for multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by the their values defined in the first programming set.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

## Parameters

*\*prog*

Specifies the events and the modes to use in the Performance Monitor setup. The *prog* parameter supports the following modes:

### **PM\_USER**

Counts processes running in the User Mode (default is set to Off).

### **PM\_KERNEL**

Counts processes running in the Kernel Mode (default is set to Off).

### **PM\_COUNT**

Starts counting immediately (default is set to Not to Start Counting).

### **PM\_PROCESS**

Creates a process-level counting group.

*\*prog\_mm*

Specifies the events and the modes to use in the Performance Monitor setup. The *prog\_mm* parameter supports the following modes:

### **PM\_USER**

Counts processes running in the User Mode (default is set to Off).

### **PM\_KERNEL**

Counts processes running in the Kernel Mode (default is set to Off).

### **PM\_COUNT**

Starts counting immediately (default is set to Not to start counting).

### **PM\_PROCTREE**

Sets counting to On only for the calling process and its descendants (default is set to Off).

The *PM\_PROCTREE* mode and the *PM\_COUNT* mode defined in the first setting fix the value for the counting.

## Return Values

0 Operation completed successfully.

Positive Error Code                      Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h`                      Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_mythread\_mx and pm\_get\_program\_mythread\_mm Subroutines” on page 1072, “pm\_delete\_program\_mythread Subroutine” on page 1031, “pm\_get\_data\_mythread\_mx or pm\_get\_tdata\_mythread\_mx Subroutine” on page 1049, “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_stop\_mythread and pm\_tstop\_mythread Subroutine ” on page 1133, “pm\_reset\_data\_mythread Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_pgroup Subroutine

### Purpose

Sets Performance Monitor programming for a target pthread and creates a counting group.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_pgroup ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_t *prog;
```

### Description

The **pm\_set\_program\_pgroup** subroutine sets the Performance Monitor programming for a target pthread. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process. The setting includes the events to be counted and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

This call also creates a counting group, which includes the target pthread and any pthread that it, or any of its descendants, will create in the future. Optionally, the group can be defined as also containing all the existing and future pthreads belonging to the target process.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The counting mode includes User Mode and/or Kernel Mode, and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_pgroup** subroutine is called.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Specifies the event modes to use in Performance Monitor setup. The following modes are supported: <b>PM_USER</b> Counts processes running in User Mode (default is set to Off) <b>PM_KERNEL</b> Counts processes running in Kernel Mode (default is set to Off) <b>PM_COUNT</b> Starts counting immediately (default is set to Not to Start Counting) <b>PM_PROCESS</b> Creates a process-level counting group

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<code>/usr/include/pmapi.h</code>	Defines standard macros, data types, and subroutines.
-----------------------------------	---

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_get\_program\_pgroup Subroutine” on page 1074, “pm\_initialize Subroutine” on page 1085,

“pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_pgroup\_mx and pm\_set\_program\_pgroup\_mm Subroutines

### Purpose

Sets Performance Monitor programming in counter multiplexing mode and multi-mode for a target pthread and creates a counting group.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_pgroup_mx ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_mx_t *prog;
```

```
int pm_set_program_pgroup_mm ( pid, tid, ptid, *prog_mm)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_mm_t *prog_mm;
```

### Description

The **pm\_set\_program\_pgroup\_mx** and the **pm\_set\_program\_pgroup\_mm** subroutines set the Performance Monitor programming respectively in counter multiplexing mode or in multi-mode for a target pthread. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process.

The **pm\_set\_program\_pgroup\_mx** setting includes the list of the event arrays to be counted and a mode in which to count. The mode is global to all of the event lists. The events to count are in an array of list of event identifiers.

The **pm\_set\_program\_pgroup\_mm** setting includes the lists of the event arrays to be counted and the associated mode in which to count each event array. A counting mode is defined for each event array.

The event identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

Both subroutines create a counting group, which includes the target pthread and any pthread that it, or any of its descendants, will create in the future. Optionally, the group can be defined as also containing all the existing and future pthreads belonging to the target process.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The counting mode for both subroutines includes the User Mode, or the Kernel Mode, or both of them; and the Initial Counting State. The defaults are set to Off for the User Mode and the Kernel Mode, and the initial default state is set to delay counting until the **pm\_start\_pgroup** subroutine is called.

When you use the **pm\_set\_program\_pgroup\_mm** subroutine for multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by their values defined in the first programming set.

If the list includes an event that can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Specifies the events and the modes to use in the Performance Monitor setup. The <i>prog</i> parameter supports the following modes: <b>PM_USER</b> Counts processes running in the User Mode (default is set to Off). <b>PM_KERNEL</b> Counts processes running in the Kernel Mode (default is set to Off). <b>PM_COUNT</b> Starts counting immediately (default is set to Not to Start Counting). <b>PM_PROCESS</b> Creates a process-level counting group.



*\*prog\_mm*

Specifies the events and the modes to use in the Performance Monitor setup. The *prog\_mm* parameter supports the following modes:

**PM\_USER**

Counts processes running in the User Mode (default is set to Off).

**PM\_KERNEL**

Counts processes running in the Kernel Mode (default is set to Off).

**PM\_COUNT**

Starts counting immediately (default is set to Not to start counting).

**PM\_PROCTREE**

Sets counting to On only for the calling process and its descendants (default is set to Off).

The *PM\_PROCTREE* mode and the *PM\_COUNT* mode defined in the first setting fix the value for the counting.

## Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup\_mx and pm\_get\_tdata\_pgroup\_mx Subroutine” on page 1052, “pm\_get\_program\_pgroup\_mx and pm\_get\_program\_pgroup\_mm Subroutines” on page 1075, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126, “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_pthread Subroutine

### Purpose

Sets Performance Monitor programming for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_thread ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_t *prog;
```

## Description

The **pm\_set\_program\_thread** subroutine sets the Performance Monitor programming for a target pthread. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process. The setting includes the events to be counted and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The counting mode includes User Mode and/or Kernel Mode, and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the Initial Default State is set to delay counting until the **pm\_start\_thread** subroutine is called.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Specifies the event modes to use in Performance Monitor setup. The following modes are supported: <b>PM_USER</b> Counts processes running in User Mode (default is set to Off) <b>PM_KERNEL</b> Counts processes running in Kernel Mode (default is set to Off) <b>PM_COUNT</b> Starts counting immediately (default is set to Not to Start Counting)

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the "pm_error Subroutine" on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_thread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1054, “pm\_get\_program\_thread Subroutine” on page 1077, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_thread Subroutine” on page 1093, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1127, “pm\_stop\_thread and pm\_tstop\_thread Subroutine ” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_thread\_mx and pm\_set\_program\_thread\_mm Subroutines

### Purpose

Sets Performance Monitor programming in counter multiplexing mode and multi-mode for a target thread.

### Library

Performance Monitor APIs Library (`libpmapi.a`)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_thread_mx ( pid, tid, ptid, *prog)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_mx_t *prog;
```

```
int pm_set_program_thread_mm ( pid, tid, ptid, *prog_mm)
pid_t pid;
tid_t tid;
ptid_t ptid;
pm_prog_mm_t *prog_mm;
```

### Description

The `pm_set_program_thread_mx` and the `pm_set_program_thread_mm` subroutines set the Performance Monitor programming respectively in counter multiplexing mode or in multi-mode for a target thread. The thread must be stopped and must be part of a debuggee process, under the control of the calling process.

The `pm_set_program_thread_mx` setting includes the list of the event arrays events to be counted and a mode in which to count. The mode is global to all of the event lists. The events to count are in an array of list of event identifiers.

The **pm\_set\_program\_thread\_mm** subroutine setting includes the list of the event arrays to be counted, and the associated mode in which to count each event array. A counting mode is defined for each event array.

The event identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

The counting mode for both subroutines includes the User Mode or the Kernel Mode, or both; and the Initial Counting State. The defaults are set to Off for the User Mode and the Kernel Mode, and the Initial Default State is set to delay counting until the **pm\_start\_thread** subroutine is called.

When you use the **pm\_set\_program\_thread\_mm** subroutine for multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by their values defined in the first programming set.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_initialize** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*prog</i>	Specifies the events and the modes to use in the Performance Monitor setup. The <i>prog</i> parameter supports the following modes: <b>PM_USER</b> Counts processes running in the User Mode (default is set to Off). <b>PM_KERNEL</b> Counts processes running in the Kernel Mode (default is set to Off). <b>PM_COUNT</b> Starts counting immediately (default is set to Not to Start Counting).

*\*prog\_mm*

Specifies the events and the associated modes to use in the Performance Monitor setup. The *prog\_mm* parameter supports the following modes:

**PM\_USER**

Counts processes running in the User Mode (default is set to Off).

**PM\_KERNEL**

Counts processes running in the Kernel Mode (default is set to Off).

**PM\_COUNT**

Starts counting immediately (default is set to Not to start counting).

**PM\_PROCTREE**

Sets counting to On only for the calling process and its descendants (default is set to Off).

The *PM\_PROCTREE* mode and the *PM\_COUNT* mode defined in the first setting fix the value for the counting.

## Return Values

**0** Operation completed successfully.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_thread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_thread\_mx or pm\_get\_tdata\_thread\_mx Subroutine” on page 1055, “pm\_get\_program\_thread\_mx and pm\_get\_program\_thread\_mm Subroutines” on page 1078, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_thread Subroutine” on page 1093, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1127, “pm\_stop\_thread and pm\_tstop\_thread Subroutine ” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_thread Subroutine

### Purpose

Sets Performance Monitor programming for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_thread ( pid, tid, *prog)
pid_t pid;
tid_t tid;
pm_prog_t *prog;
```

## Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_set\_program\_pthread** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_set\_program\_pthread** subroutine with a *ptid* parameter equal to 0.

The **pm\_set\_program\_thread** subroutine sets the Performance Monitor programming for a target kernel thread. The thread must be stopped and must be part of a debuggee process, under the control of the calling process. The setting includes the events to be counted and a mode in which to count. The events to count are in a list of event identifiers. The identifiers must be selected from the lists returned by the **pm\_init** subroutine.

The counting mode includes User Mode and/or Kernel Mode, and the Initial Counting State. The defaults are set to Off for User Mode and Kernel Mode, and the Initial Default State is set to delay counting until the **pm\_start\_thread** subroutine is called.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Process ID of target thread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target thread.
<i>*prog</i>	Specifies the event modes to use in Performance Monitor setup. The following modes are supported: <b>PM_USER</b> Counts processes running in User Mode (default is set to Off) <b>PM_KERNEL</b> Counts processes running in Kernel Mode (default is set to Off) <b>PM_COUNT</b> Starts counting immediately (default is set to Not to Start Counting)

## Return Values

0	Operation completed successfully.
Positive Error Code	Refer to the <b>pm_error</b> ("pm_error Subroutine" on page 1035) subroutine to decode the error code.

## Error Codes

Refer to the **pm\_error** ("pm\_error Subroutine" on page 1035) subroutine.

## Files

`/usr/include/pmapi.h`

Defines standard macros, data types, and subroutines.

## Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_get\_program\_thread** (“pm\_get\_program\_thread Subroutine” on page 1080) subroutine, **pm\_delete\_program\_thread** (“pm\_delete\_program\_thread Subroutine” on page 1034) subroutine, **pm\_get\_data\_thread** (“pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1057) subroutine, **pm\_start\_thread** (“pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1129) subroutine, **pm\_stop\_thread** (“pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137) subroutine, **pm\_reset\_data\_thread** (“pm\_reset\_data\_thread Subroutine” on page 1094) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_set\_program\_thread\_mx and pm\_set\_program\_thread\_mm Subroutines

### Purpose

Sets Performance Monitor programming in counter multiplexing mode and multi-mode for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_set_program_thread_mx ( pid, tid, *prog)  
pid_t pid;  
tid_t tid;  
pm_prog_mx_t *prog;
```

```
int pm_set_program_thread_mm ( pid, tid, *prog_mm)  
pid_t pid;  
tid_t tid;  
pm_prog_mm_t *prog_mm;
```

### Description

The **pm\_set\_program\_thread\_mx** and the **pm\_set\_program\_thread\_mm** subroutines support only the 1:1 threading model. They have been superseded respectively by the **pm\_set\_program\_pthread\_mx** and the **pm\_set\_program\_pthread\_mm** subroutines, which support both the 1:1 and the M:N threading models. A call to the **pm\_set\_program\_thread\_mx** subroutine or the **pm\_set\_program\_thread\_mm** subroutine is respectively equivalent to a call to the **pm\_set\_program\_pthread\_mx** subroutine or the **pm\_set\_program\_pthread\_mm** subroutine with a *ptid* parameter equal to 0.

The **pm\_set\_program\_thread\_mx** and the **pm\_set\_program\_thread\_mm** subroutines set the Performance Monitor programming respectively in counter multiplexing mode or multi-mode for a target kernel thread. The thread must be stopped and must be part of a debuggee process, under the control of the calling process.

The **pm\_set\_program\_thread\_mx** setting includes the list of the event arrays to be counted and a mode in which to count. The mode is global to all of the event lists. The events to count are in an array of list of event identifiers.

The **pm\_set\_program\_thread\_mm** setting includes the list of the event arrays to be counted, and the associated mode in which to count each event array. A counting mode is defined for each event array.

The event identifiers must be selected from the lists returned by the **pm\_initialize** subroutine.

The counting mode for both subroutines includes the User Mode, or the Kernel Mode, or both of them; and the Initial Counting State. The defaults are set to Off for the User Mode and the Kernel Mode, and the Initial Default State is set to delay counting until the **pm\_start\_thread** subroutine is called.

When you use the **pm\_set\_program\_thread\_mm** subroutine for the multi-mode counting, the Process Tree Mode and the Start Counting Mode are fixed by their values in the first programming set.

If the list includes an event which can be used with a threshold (as indicated by the **pm\_init** subroutine), a threshold value can also be specified.

## Parameters

<i>pid</i>	Process ID of target thread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target thread.
<i>*prog</i>	Specifies the events and the modes to use in the Performance Monitor setup. The <i>prog</i> parameter supports the following modes: <b>PM_USER</b> Counts processes running in the User Mode (default is set to Off). <b>PM_KERNEL</b> Counts processes running in the Kernel Mode (default is set to Off). <b>PM_COUNT</b> Starts counting immediately (default is set to Not to Start Counting).
<i>*prog_mm</i>	Specifies the events and the associated modes to use in the Performance Monitor setup. The <i>prog_mm</i> parameter supports the following modes: <b>PM_USER</b> Counts processes running in the User Mode (default is set to Off). <b>PM_KERNEL</b> Counts processes running in the Kernel Mode (default is set to Off). <b>PM_COUNT</b> Starts counting immediately (default is set to Not to start counting). <b>PM_PROCTREE</b> Sets counting to On only for the calling process and its descendants (default is set to Off).  The <i>PM_PROCTREE</i> mode and the <i>PM_COUNT</i> mode defined in the first setting fix the value for the counting.



## Return Values

0 Operation completed successfully.  
Positive error code Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_get\_program\_thread\_mx and pm\_get\_program\_thread\_mm Subroutines” on page 1082, “pm\_delete\_program\_thread Subroutine” on page 1034, “pm\_get\_data\_thread\_mx or pm\_get\_tdata\_thread\_mx Subroutine” on page 1059, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1129, “pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137, “pm\_reset\_data\_thread Subroutine” on page 1094.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_start and pm\_tstart Subroutine

### Purpose

Starts system wide Performance Monitor counting.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_start()

int pm_tstart(*time)
timebasestruct_t *time;
```

### Description

The **pm\_start** subroutine starts system wide Performance Monitor counting.

The **pm\_tstart** subroutine starts system wide Performance Monitor counting, and returns a timestamp indicating when the counting was started.

### Parameters

*\*time* Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

0 Operation completed successfully.  
Positive error code Refer to the “pm\_error Subroutine” on page 1035 to decode the error code

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program Subroutine” on page 1095, “pm\_get\_program Subroutine” on page 1061, “pm\_delete\_program Subroutine” on page 1028, “pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine” on page 1036, “pm\_stop and pm\_tstop Subroutine ” on page 1130, “pm\_reset\_data Subroutine” on page 1087.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_start\_group and pm\_tstart\_group Subroutine

### Purpose

Starts Performance Monitor counting for the counting group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_start_group ( pid, tid)
pid_t pid;
tid_t tid;
```

```
int pm_tstart_group ( pid, tid, *time)
pid_t pid;
tid_t tid;
timebasestruct_t *time
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_start\_pgroup** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_start\_pgroup** subroutine with a *ptid* parameter equal to 0.

The **pm\_start\_group** subroutine starts the Performance Monitor counting for a target kernel thread and the counting group to which it belongs. This counting is effective immediately for the target thread. For all the other thread members of the counting group, the counting will start after their next redispach, but only

if their current counting state is already set to On. The counting state of a thread in a group is obtained by ANDing the thread counting state with the group state. If their counting state is currently set to Off, no counting starts until they call either the **pm\_start\_mythread** subroutine or the **pm\_start\_mygroup** themselves, or until a debugger process calls the **pm\_start\_thread** subroutine or the **pm\_start\_group** subroutine on their behalf.

The **pm\_tstart\_group** subroutine starts the Performance Monitor counting for a target kernel thread and the counting group to which it belongs, and returns a timestamp indicating when the counting was started.

## Parameters

<i>pid</i>	Process ID of target thread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target thread.
<i>*time</i>	Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<i>/usr/include/pmapi.h</i>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_group Subroutine” on page 1096, “pm\_get\_program\_group Subroutine” on page 1062, “pm\_delete\_program\_group Subroutine” on page 1029, “pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine” on page 1039, “pm\_stop\_group and pm\_tstop\_group Subroutine” on page 1131, “pm\_reset\_data\_group Subroutine” on page 1088.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_start\_mygroup and pm\_tstart\_mygroup Subroutine

### Purpose

Starts Performance Monitor counting for the group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

## Syntax

```
#include <pmapi.h>
```

```
int pm_start_mygroup()
```

```
int pm_tstart_mygroup (*time)  
timebasestruct_t *time
```

## Description

The **pm\_start\_mygroup** subroutine starts the Performance Monitor counting for the calling kernel thread and the counting group to which it belongs. Counting is effective immediately for the calling thread. For all the other threads members of the counting group, the counting starts after their next redispach, but only if their current counting state is already set to On. The counting state of a thread in a group is obtained by ANDing the thread counting state with the group state. If their counting state is currently set to Off, no counting starts until they call either the **pm\_start\_mythread** subroutine or the **pm\_start\_mygroup** subroutine themselves, or until a debugger process calls the **pm\_start\_thread** subroutine or the **pm\_start\_group** subroutine on their behalf.

The **pm\_tstart\_mygroup** subroutine starts the Performance Monitor counting for the calling kernel thread and the counting group to which it belongs, and returns a timestamp indicating when the counting was started.

## Parameters

*\*time*                      Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

0                              Operation completed successfully.  
Positive error code            Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h**                      Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mygroup Subroutine” on page 1102, “pm\_get\_program\_mygroup Subroutine” on page 1068, “pm\_delete\_program\_mygroup Subroutine” on page 1030, “pm\_get\_data\_mygroup, pm\_get\_tdata\_mygroup or pm\_get\_Tdata\_mygroup Subroutine” on page 1044, “pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine ” on page 1132, “pm\_reset\_data\_mygroup Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_start\_mythread and pm\_tstart\_mythread Subroutine

### Purpose

Starts Performance Monitor counting for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_start_mythread()
```

```
int pm_tstart_mythread(*time)  
timebasestruct_t *time;
```

### Description

The **pm\_start\_mythread** subroutine starts Performance Monitor counting for the calling kernel thread. Counting is effective immediately unless the thread is in a group, and that group's counting is not currently set to On. The counting state of a thread in a group is obtained by ANDing the thread counting state with the group state.

The **pm\_tstart\_mythread** subroutine starts Performance Monitor counting for the calling kernel thread, and returns a timestamp indicating when the counting was started.

### Parameters

*\*time*                      Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the **time\_base\_to\_time** subroutine.

### Return Values

0                              Operation completed successfully.  
Positive Error Code           Refer to the **pm\_error** ("pm\_error Subroutine" on page 1035) subroutine to decode the error code.

### Error Codes

Refer to the **pm\_error** ("pm\_error Subroutine" on page 1035) subroutine

### Files

**/usr/include/pmapi.h**                      Defines standard macros, data types, and subroutines.

### Related Information

The **pm\_init** ("pm\_init Subroutine" on page 1084) subroutine, **pm\_error** ("pm\_error Subroutine" on page 1035) subroutine, **pm\_set\_program\_mythread** ("pm\_set\_program\_mythread Subroutine" on page 1106) subroutine, **pm\_get\_program\_mythread** ("pm\_get\_program\_mythread Subroutine" on page 1071) subroutine, **pm\_delete\_program\_mythread** ("pm\_delete\_program\_mythread Subroutine" on page 1031) subroutine, **pm\_get\_data\_mythread** ("pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine" on page 1047) subroutine, **pm\_stop\_mythread** ("pm\_stop\_mythread

and `pm_tstop_mythread` Subroutine ” on page 1133) subroutine, `pm_reset_data_mythread` (“`pm_reset_data_mythread` Subroutine” on page 1090) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## `pm_start_pgroup` and `pm_tstart_pgroup` Subroutine

### Purpose

Starts Performance Monitor counting for the counting group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_start_pgroup ( pid, tid, ptid)
pid_t pid;
tid_t tid;
ptid_t ptid;
```

```
int pm_tstart_pgroup ( pid, tid, ptid, *time)
pid_t pid;
tid_t tid;
ptid_t ptid;
timebasestruct_t *time
```

### Description

The `pm_start_pgroup` subroutine starts the Performance Monitor counting for a target pthread and the counting group to which it belongs. This counting is effective immediately for the target pthread. For all the other thread members of the counting group, the counting will start after their next redispatch, but only if their current counting state is already set to On. The counting state of a pthread in a group is obtained by ANDing the pthread counting state with the group state. If their counting state is currently set to Off, no counting starts until they call either the `pm_start_mythread` subroutine or the `pm_start_mygroup` themselves, or until a debugger process calls the `pm_start_pthread` subroutine or the `pm_start_pgroup` subroutine on their behalf.

The `pm_tstart_pgroup` subroutine starts the Performance Monitor counting for a target pthread and the counting group to which it belongs, and returns a timestamp indicating when the counting was started.

If the pthread is running in 1:1 mode, only the `tid` parameter must be specified. If the pthread is running in m:n mode, only the `ptid` parameter must be specified. If both the `ptid` and `tid` parameters are specified, they must be referring to a single pthread with the `ptid` parameter specified and currently running on a kernel thread with specified `tid` parameter.

### Parameters

<code>pid</code>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<code>tid</code>	Thread ID of target pthread. To ignore this parameter, set it to 0.

*ptid*

Pthread ID of the target pthread. To ignore this parameter, set it to 0.

*\*time*

Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

0

Operation completed successfully.

Positive error code

Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

*/usr/include/pmapi.h*

Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_get\_program\_pgroup Subroutine” on page 1074, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_set\_program\_pgroup Subroutine” on page 1109, “pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine ” on page 1134.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_start\_pthread and pm\_tstart\_pthread Subroutine

### Purpose

Starts Performance Monitor counting for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_start_pthread ( pid, tid, ptid)
```

```
pid_t pid;
```

```
tid_t tid;
```

```
ptid_t ptid;
```

```
int pm_start_pthread ( pid, tid, ptid, *time)
```

```
pid_t pid;
```

```
tid_t tid;
```

```
ptid_t ptid;
```

```
timebasestruct_t *time
```

## Description

The **pm\_start\_thread** subroutine starts Performance Monitor counting for a target pthread. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process. Counting is effective immediately unless the thread is in a group and the group counting is not currently set to On. The counting state of a thread in a group is obtained by ANDing the thread counting state with the group state.

The **pm\_tstart\_thread** subroutine starts Performance Monitor counting for a target pthread, and returns a timestamp indicating when the counting was started.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*time</i>	Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<b>/usr/include/pmapi.h</b>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_delete\_program\_thread Subroutine” on page 1033, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1054, “pm\_get\_program\_thread Subroutine” on page 1077, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_thread Subroutine” on page 1093, “pm\_set\_program\_thread Subroutine” on page 1113, “pm\_stop\_thread and pm\_tstop\_thread Subroutine ” on page 1136.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.



---

## pm\_start\_thread and pm\_tstart\_thread Subroutine

### Purpose

Starts Performance Monitor counting for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_start_thread (pid, tid)
```

```
pid_t pid;
```

```
tid_t tid;
```

```
int pm_tstart_thread ( pid, tid, *time)
```

```
pid_t pid;
```

```
tid_t tid;
```

```
timebasestruct_t *time
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_start\_pthread** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_start\_pthread** subroutine with a *ptid* parameter equal to 0.

The **pm\_start\_thread** subroutine starts Performance Monitor counting for a target kernel thread. The thread must be stopped and must be part of a debuggee process, under the control of the calling process. Counting is effective immediately unless the thread is in a group and the group counting is not currently set to On. The counting state of a thread in a group is obtained by ANDing the thread counting state with the group state.

The **pm\_tstart\_thread** subroutine starts Performance Monitor counting for a target kernel thread, and returns a timestamp indicating when the counting was started.

### Parameters

*pid*

Process ID of target thread. Target process must be a debuggee of the caller process.

*tid*

Thread ID of target thread.

*\*time*

Pointer to a structure containing the timebase value when the counting was started. This can be converted to time using the **time\_base\_to\_time** subroutine.

### Return Values

0

Operation completed successfully.

Positive Error Code

Refer to the **pm\_error** ("pm\_error Subroutine" on page 1035) subroutine to decode the error code.

### Error Codes

Refer to the **pm\_error** ("pm\_error Subroutine" on page 1035) subroutine.

## Files

`/usr/include/pmapi.h` Defines standard macros, data types, and subroutines.

## Related Information

The **pm\_init** (“pm\_init Subroutine” on page 1084) subroutine, **pm\_error** (“pm\_error Subroutine” on page 1035) subroutine, **pm\_set\_program\_thread** (“pm\_set\_program\_thread Subroutine” on page 1117) subroutine, **pm\_get\_program\_thread** (“pm\_get\_program\_thread Subroutine” on page 1080) subroutine, **pm\_delete\_program\_thread** (“pm\_delete\_program\_thread Subroutine” on page 1034) subroutine, **pm\_get\_data\_thread** (“pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1057) subroutine, **pm\_stop\_thread** (“pm\_stop\_thread and pm\_tstop\_thread Subroutine” on page 1137) subroutine, **pm\_reset\_data\_thread** (“pm\_reset\_data\_thread Subroutine” on page 1094) subroutine.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_stop and pm\_tstop Subroutine

### Purpose

Stops system wide Performance Monitor counting.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_stop ()

int pm_tstop(*time)
timebasestruct_t *time;
```

### Description

The **pm\_stop** subroutine stops system wide Performance Monitoring counting.

The **pm\_tstop** subroutine stops system wide Performance Monitoring counting, and returns a timestamp indicating when the counting was stopped.

### Parameters

*\*time* Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the **time\_base\_to\_time** subroutine.

### Return Values

**0** Operation completed successfully.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

### Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

`/usr/include/pmapi.h`

Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program Subroutine” on page 1095, “pm\_get\_program Subroutine” on page 1061, “pm\_delete\_program Subroutine” on page 1028, “pm\_get\_data, pm\_get\_tdata, pm\_get\_Tdata, pm\_get\_data\_cpu, pm\_get\_tdata\_cpu, pm\_get\_Tdata\_cpu, pm\_get\_data\_lcpu, pm\_get\_tdata\_lcpu and pm\_get\_Tdata\_lcpu Subroutine” on page 1036, “pm\_start and pm\_tstart Subroutine” on page 1121, “pm\_reset\_data Subroutine” on page 1087.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_stop\_group and pm\_tstop\_group Subroutine

### Purpose

Stops Performance Monitor counting for the group to which a target thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_stop_group ( pid, tid )
pid_t pid;
tid_t tid;
```

```
int pm_tstop_group ( pid, tid, *time )
pid_t pid;
tid_t tid;
timebasestruct_t *time;
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **pm\_stop\_pgroup** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **pm\_stop\_pgroup** subroutine with a *ptid* parameter equal to 0.

The **pm\_stop\_group** subroutine stops Performance Monitor counting for a target kernel thread, the counting group to which it belongs, and all the other thread members of the same group. Counting stops immediately for all the threads in the counting group. The target thread must be stopped and must be part of a debuggee process, under control of the calling process.

The **pm\_tstop\_group** subroutine stops Performance Monitor counting for a target kernel thread, the counting group to which it belongs, and all the other thread members of the same group, and returns a timestamp indicating when the counting was stopped.

### Parameters

*pid*

Process ID of target thread. Target process must be a debuggee of the caller process.

*tid*  
*\*time*

Thread ID of target thread.  
Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

**0** Operation completed successfully.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_group Subroutine” on page 1096, “pm\_get\_program\_group Subroutine” on page 1062, “pm\_delete\_program\_group Subroutine” on page 1029, “pm\_get\_data\_group, pm\_get\_tdata\_group and pm\_get\_Tdata\_group Subroutine” on page 1039, “Syntax” on page 1122, “pm\_reset\_data\_group Subroutine” on page 1088.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_stop\_mygroup and pm\_tstop\_mygroup Subroutine

### Purpose

Stops Performance Monitor counting for the group to which the calling thread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>

int pm_stop_mygroup ()

int pm_tstop_mygroup(*time)
timebasestruct_t *time;
```

### Description

The **pm\_stop\_mygroup** subroutine stops Performance Monitor counting for the group to which the calling kernel thread belongs. This is effective immediately for all the threads in the counting group.

The **pm\_tstop\_mygroup** subroutine stops Performance Monitor counting for the group to which the calling kernel thread belongs, and returns a timestamp indicating when the counting was stopped.

## Parameters

*\*time* Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

**0** Operation completed successfully.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mygroup Subroutine” on page 1102, “pm\_get\_program\_mygroup Subroutine” on page 1068, “pm\_delete\_program\_mygroup Subroutine” on page 1030, “pm\_get\_data\_mygroup, pm\_get\_tdata\_mygroup or pm\_get\_Tdata\_mygroup Subroutine” on page 1044, “Description” on page 1124, “pm\_reset\_data\_mygroup Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_stop\_mythread and pm\_tstop\_mythread Subroutine

### Purpose

Stops Performance Monitor counting for the calling thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_stop_mythread ()
```

```
int pm_tstop_mythread(*time)  
timebasestruct_t *time;
```

### Description

The **pm\_stop\_mythread** subroutine stops Performance Monitor counting for the calling kernel thread.

The **pm\_tstop\_mythread** subroutine stops Performance Monitor counting for the calling kernel thread, and returns a timestamp indicating when the counting was stopped.

## Parameters

*\*time* Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

**0** Operation completed successfully.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_mythread Subroutine” on page 1106, “pm\_get\_program\_mythread Subroutine” on page 1071, “pm\_delete\_program\_mythread Subroutine” on page 1031, “pm\_get\_data\_mythread, pm\_get\_tdata\_mythread or pm\_get\_Tdata\_mythread Subroutine” on page 1047, “pm\_start\_mythread and pm\_tstart\_mythread Subroutine” on page 1125, “pm\_reset\_data\_mythread Subroutine” on page 1090.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_stop\_pgroup and pm\_tstop\_pgroup Subroutine

### Purpose

Stops Performance Monitor counting for the group to which a target pthread belongs.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_stop_pgroup ( pid, tid, ptid)
pid_t pid;
tid_t tid;
ptid_t ptid;
```

```
int pm_tstop_pgroup ( pid, tid, ptid, *time)
pid_t pid;
tid_t tid;
ptid_t ptid;
timebasestruct_t *time;
```

## Description

The **pm\_stop\_pgroup** subroutine stops Performance Monitor counting for a target pthread, the counting group to which it belongs, and all the other pthread members of the same group. Counting stops immediately for all the pthreads in the counting group. The target pthread must be stopped and must be part of a debuggee process, under control of the calling process.

The **pm\_tstop\_pgroup** subroutine stops Performance Monitor counting for a target pthread, the counting group to which it belongs, and all the other pthread members of the same group, and returns a timestamp indicating when the counting was stopped.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

## Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*time</i>	Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the <b>time_base_to_time</b> subroutine.

## Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

<b>/usr/include/pmapi.h</b>	Defines standard macros, data types, and subroutines.
-----------------------------	---

## Related Information

The “pm\_delete\_program\_pgroup Subroutine” on page 1032, “pm\_error Subroutine” on page 1035, “pm\_get\_data\_pgroup, pm\_get\_tdata\_pgroup and pm\_get\_Tdata\_pgroup Subroutine” on page 1050, “pm\_get\_program\_pgroup Subroutine” on page 1074, “pm\_initialize Subroutine” on page 1085, “pm\_reset\_data\_pgroup Subroutine” on page 1091, “pm\_set\_program\_pgroup Subroutine” on page 1109, “pm\_start\_pgroup and pm\_tstart\_pgroup Subroutine” on page 1126.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## pm\_stop\_pthread and pm\_tstop\_pthread Subroutine

### Purpose

Stops Performance Monitor counting for a target pthread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_stop_pthread ( pid, tid, ptid)
```

```
pid_t pid;
```

```
tid_t tid;
```

```
ptid_t ptid;
```

```
int pm_tstop_pthread ( pid, tid, ptid, *time)
```

```
pid_t pid;
```

```
tid_t tid;
```

```
ptid_t ptid;
```

```
timebasestruct_t *time;
```

### Description

The **pm\_stop\_pthread** subroutine stops Performance Monitor counting for a target pthread. The pthread must be stopped and must be part of a debuggee process, under the control of the calling process.

The **pm\_tstop\_pthread** subroutine stops Performance Monitor counting for a target pthread, and returns a timestamp indicating when the counting was stopped.

If the pthread is running in 1:1 mode, only the *tid* parameter must be specified. If the pthread is running in m:n mode, only the *ptid* parameter must be specified. If both the *ptid* and *tid* parameters are specified, they must be referring to a single pthread with the *ptid* parameter specified and currently running on a kernel thread with specified *tid* parameter.

### Parameters

<i>pid</i>	Process ID of target pthread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target pthread. To ignore this parameter, set it to 0.
<i>ptid</i>	Pthread ID of the target pthread. To ignore this parameter, set it to 0.
<i>*time</i>	Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the <b>time_base_to_time</b> subroutine.

### Return Values

0	Operation completed successfully.
Positive error code	Refer to the “pm_error Subroutine” on page 1035 to decode the error code.

### Error Codes

Refer to the “pm\_error Subroutine” on page 1035.



## Files

`/usr/include/pmapi.h`

Defines standard macros, data types, and subroutines.

## Related Information

The “`pm_delete_program_thread` Subroutine” on page 1033, “`pm_error` Subroutine” on page 1035, “`pm_get_data_thread`, `pm_get_tdata_thread` or `pm_get_Tdata_thread` Subroutine” on page 1054, “`pm_get_program_thread` Subroutine” on page 1077, “`pm_initialize` Subroutine” on page 1085, “`pm_reset_data_thread` Subroutine” on page 1093, “`pm_set_program_thread` Subroutine” on page 1113, “`pm_start_thread` and `pm_tstart_thread` Subroutine” on page 1127.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## `pm_stop_thread` and `pm_tstop_thread` Subroutine

### Purpose

Stops Performance Monitor counting for a target thread.

### Library

Performance Monitor APIs Library (**libpmapi.a**)

### Syntax

```
#include <pmapi.h>
```

```
int pm_stop_thread (pid, tid)
pid_t pid;
tid_t tid;
```

```
int pm_tstop_thread (pid, tid, *time)
pid_t pid;
tid_t tid;
timebasestruct_t *time;
```

### Description

This subroutine supports only the 1:1 threading model. It has been superseded by the **`pm_stop_thread`** subroutine, which supports both the 1:1 and the M:N threading models. A call to this subroutine is equivalent to a call to the **`pm_stop_thread`** subroutine with a *ptid* parameter equal to 0.

The **`pm_stop_thread`** subroutine stops Performance Monitor counting for a target kernel thread. The thread must be stopped and must be part of a debuggee process, under the control of the calling process.

The **`pm_tstop_thread`** subroutine stops Performance Monitor counting for a target kernel thread, and returns a timestamp indicating when the counting was stopped.

### Parameters

<i>pid</i>	Process ID of target thread. Target process must be a debuggee of the caller process.
<i>tid</i>	Thread ID of target thread.

*\*time*

Pointer to a structure containing the timebase value when the counting was stopped. This can be converted to time using the **time\_base\_to\_time** subroutine.

## Return Values

**0** Operation completed successfully.  
**Positive error code** Refer to the “pm\_error Subroutine” on page 1035 to decode the error code.

## Error Codes

Refer to the “pm\_error Subroutine” on page 1035.

## Files

**/usr/include/pmapi.h** Defines standard macros, data types, and subroutines.

## Related Information

The “pm\_init Subroutine” on page 1084, “pm\_error Subroutine” on page 1035, “pm\_set\_program\_thread Subroutine” on page 1117, “pm\_get\_program\_thread Subroutine” on page 1080, “pm\_delete\_program\_thread Subroutine” on page 1034, “pm\_get\_data\_thread, pm\_get\_tdata\_thread or pm\_get\_Tdata\_thread Subroutine” on page 1057, “pm\_start\_thread and pm\_tstart\_thread Subroutine” on page 1129, “pm\_reset\_data\_thread Subroutine” on page 1094.

Performance Monitor API Programming Concepts in *AIX 5L Version 5.3 Performance Tools Guide and Reference*.

---

## poll Subroutine

### Purpose

Checks the I/O status of multiple file descriptors and message queues.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/poll.h>
#include <sys/select.h>
#include <sys/types.h>
```

```
int poll( ListPointer, Nfdmsgs, Timeout)
void *ListPointer;
unsigned long Nfdmsgs;
long Timeout;
```

### Description

The **poll** subroutine checks the specified file descriptors and message queues to see if they are ready for reading (receiving) or writing (sending), or to see if they have an exceptional condition pending. Even though there are **OPEN\_MAX** number of file descriptors available, only **FD\_SETSIZE** number of file descriptors are accessible with this subroutine.

**Note:** The **poll** subroutine applies only to character devices, pipes, message queues, and sockets. Not all character device drivers support it. See the descriptions of individual character devices for information about whether and how specific device drivers support the **poll** and **select** subroutines.

For compatibility with previous releases of this operating system and with BSD systems, the **select** subroutine is also supported.

In AIX 5.3 and later versions, if a program needs to use message queue support, the program source code should be compiled with the **-D\_MSGQSUPPORT** compilation flag.

## Parameters

*ListPointer* Specifies a pointer to an array of **pollfd** structures, **pollmsg** structures, or to a **pollist** structure. Each structure specifies a file descriptor or message queue ID and the events of interest for this file or message queue. The **pollfd**, **pollmsg**, and **pollist** structures are defined in the **/usr/include/sys/poll.h** file. If a **pollist** structure is to be used, a structure similar to the following should be defined in a user program. The **pollfd** structure must precede the **pollmsg** structure.

```
struct pollist {
    struct pollfd fds[3];
    struct pollmsg msgs[2];
} list;
```

The structure can then be initialized as follows:

```
list.fds[0].fd = file_descriptorA;
list.fds[0].events = requested_events;
list.msgs[0].msgid = message_id;
list.msgs[0].events = requested_events;
```

The rest of the elements in the **fds** and **msgs** arrays can be initialized the same way. The **poll** subroutine can then be called, as follows:

```
nfds = 3; /* number of pollfd structs */
nmsgs = 2; /* number of pollmsg structs */
timeout = 1000 /* number of milliseconds to timeout */
poll(&list, (nmsgs<<16)|(nfds), 1000);
```

The exact number of elements in the **fds** and **msgs** arrays must be used in the calculation of the *Nfdsmsgs* parameter.

*Nfdsmsgs* Specifies the number of file descriptors and the exact number of message queues to check. The low-order 16 bits give the number of elements in the array of **pollfd** structures, while the high-order 16 bits give the exact number of elements in the array of **pollmsg** structures. If either half of the *Nfdsmsgs* parameter is equal to a value of 0, the corresponding array is assumed not to be present.

*Timeout* Specifies the maximum length of time (in milliseconds) to wait for at least one of the specified events to occur. If the *Timeout* parameter value is -1, the **poll** subroutine does not return until at least one of the specified events has occurred. If the value of the *Timeout* parameter is 0, the **poll** subroutine does not wait for an event to occur but returns immediately, even if none of the specified events have occurred.

## poll Subroutine STREAMS Extensions

In addition to the functions described above, the **poll** subroutine multiplexes input/output over a set of file descriptors that reference open streams. The **poll** subroutine identifies those streams on which you can send or receive messages, or on which certain events occurred.

You can receive messages using the **read** subroutine or the **getmsg** system call. You can send messages using the **write** subroutine or the **putmsg** system call. Certain **streamio** operations, such as **I\_RECVFD** and **I\_SENDFD** can also be used to send and receive messages. See the **streamio** operations.

The *ListPointer* parameter specifies the file descriptors to be examined and the events of interest for each file descriptor. It points to an array having one element for each open file descriptor of interest. The array's elements are **pollfd** structures. In addition to the **pollfd** structure in the **/usr/include/sys/poll.h** file, STREAMS supports the following members:

```
int fd;           /* file descriptor */
short events;    /* requested events */
short revents;   /* returned events */
```

The **fd** field specifies an open file descriptor and the **events** and **revents** fields are bit-masks constructed by ORing any combination of the following event flags:

<b>POLLIN</b>	A nonpriority or file descriptor-passing message is present on the stream-head read queue. This flag is set even if the message is of 0 length. In the <b>revents</b> field this flag is mutually exclusive with the <b>POLLPRI</b> flag. See the <b>I_RECVFD</b> command.
<b>POLLRDNORM</b>	A nonpriority message is present on the stream-head read queue.
<b>POLLRDBAND</b>	A priority message (band > 0) is present on the stream-head read queue.
<b>POLLPRI</b>	A high-priority message is present on the stream-head read queue. This flag is set even if the message is of 0 length. In the <b>revents</b> field, this flag is mutually exclusive with the <b>POLLIN</b> flag.
<b>POLLOUT</b>	The first downstream write queue in the stream is not full. Normal priority messages can be sent at any time. See the <b>putmsg</b> system call.
<b>POLLWRNORM</b>	The same as <b>POLLOUT</b> .
<b>POLLWRBAND</b>	A priority band greater than 0 exists downstream and priority messages can be sent at anytime.
<b>POLLMSG</b>	A message containing the <b>SIGPOLL</b> signal has reached the front of the stream-head read queue.

## Return Values

On successful completion, the **poll** subroutine returns a value that indicates the total number of file descriptors and message queues that satisfy the selection criteria. The return value is similar to the *Nfdsmgs* parameter in that the low-order 16 bits give the number of file descriptors, and the high-order 16 bits give the number of message queue identifiers that had nonzero **revents** values. The **NFDS** and **NMSGs** macros, found in the **sys/select.h** file, can be used to separate these two values from the return value. The **NFDS** macro returns **NFDS#**, where the number returned indicates the number of files reporting some event or error, and the **NMSGs** macro returns **NMSGs#**, where the number returned indicates the number of message queues reporting some event or error.

A value of 0 indicates that the **poll** subroutine timed out and that none of the specified files or message queues indicated the presence of an event (all **revents** fields were values of 0).

If unsuccessful, a value of -1 is returned and the global variable **errno** is set to indicate the error.

## Error Codes

The **poll** subroutine does not run successfully if one or more of the following are true:

<b>EAGAIN</b>	Allocation of internal data structures was unsuccessful.
<b>EINTR</b>	A signal was caught during the <b>poll</b> system call and the signal handler was installed with an indication that subroutines are not to be restarted.
<b>EINVAL</b>	The number of <b>pollfd</b> structures specified by the <i>Nfdsmgs</i> parameter is greater than <b>FD_SETSIZE</b> . This error is also returned if the number of <b>pollmsg</b> structures specified by the <i>Nfdsmgs</i> parameter is greater than the maximum number of allowable message queues.
<b>EFAULT</b>	The <i>ListPointer</i> parameter in conjunction with the <i>Nfdsmgs</i> parameter addresses a location outside of the allocated address space of the process.

## Related Information

The **read** subroutine, **select** subroutine, **write** subroutine.

The **getmsg** system call, **putmsg** system call.

The **streamio** operations.

The STREAMS Overview and the Input and Output Handling Programmer's Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pollset\_create, pollset\_ctl, pollset\_destroy, pollset\_poll, and pollset\_query Subroutines

### Purpose

Check I/O status of multiple file descriptors.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/poll.h>
#include <sys/pollset.h>
#include <sys/fcntl.h>

pollset_t ps = pollset_create(int maxfd)
int rc = pollset_destroy(pollset_t ps)
int rc = pollset_ctl(pollset_t ps, struct poll_ctl *pollctl_array,
                   int array_length)
int rc = pollset_query(pollset_t ps, struct pollfd *pollfd_query)
int nfound = pollset_poll(pollset_t ps,
                         struct pollfd *polldata_array,
                         int array_length, int timeout)
```

### Description

The **pollset** application programming interface (API) efficiently poll a large file descriptor set. This interface is best used when the file descriptor set is not frequently updated. The **pollset** subroutine can provide a significant performance enhancement over traditional **select** and **poll** APIs. Improvements are most visible when the number of events returned per poll operation is small in relation to the number of file descriptors polled.

The **pollset** API uses system calls to accomplish polling. A file descriptor set (or *pollset*) is established with a successful call to **pollset\_create**. File descriptors and poll events are added, removed, or updated using the **pollset\_ctl** subroutine. The **pollset\_poll** subroutine is called to perform the poll operation. A **pollset\_query** subroutine is called to query if a file descriptor is contained in the current poll set.

A pollset is established with a successful call to **pollset\_create**. The pollset is initially empty following this system call. Each call to **pollset\_create** creates a new and independent pollset. This can be useful to applications that monitor distinct sets of file descriptors. The maximum number of file descriptors that can belong to the pollset is specified by *maxfd*. If *maxfd* has a value of -1, the maximum number of file descriptors that can belong to the pollset is bound by **OPEN\_MAX** as defined in **<sys/limits.h>** (the AIX limit of open file descriptors per process). AIX imposes a system-wide limit of 245025 active pollsets at one time. Upon failure, this system call returns -1 with **errno** set appropriately. Upon success, a pollset ID of type **pollset\_t** is returned:

```
typedef int pollset_t
```

The pollset ID must not be altered by the application. The pollset API verifies that the ID is not -1. In addition, the process ID of the application must match the process ID stored at pollset creation time.

A pollset is destroyed with a successful call to **pollset\_destroy**. Upon success, this system call returns 0. Upon failure, the **pollset\_destroy** subroutine returns -1 with **errno** set to the appropriate code. An **errno** of **EINVAL** indicates an invalid pollset ID.

File descriptors must be added to the pollset with the **pollset\_ctl** subroutine before they can be monitored. An array of **poll\_ctl** structures is passed to **pollset\_ctl** through **pollctl\_array**:

```
struct poll_ctl {
    short cmd;
    short events;
    int fd;
}
```

Each **poll\_ctl** structure contains an *fd*, *events*, and *cmd* field. The *fd* field defines the file descriptor to operate on. The *events* field contains events of interest. When *cmd* is **PS\_ADD**, the **pollset\_ctl** call adds a valid open file descriptor to the pollset. If a file descriptor is already in the pollset, **PS\_ADD** causes **pollset\_ctl** to return an error. When *cmd* is **PS\_MOD** and the file descriptor is already in the pollset, bits in the events field are added (ORed) to the monitored events. If the file descriptor is not already in the pollset, **PS\_MOD** adds a valid open file descriptor to the pollset.

Although poll events can be added by specifying an existing file descriptor, the file descriptor must be removed and then added again to remove an event. When *cmd* is **PS\_DELETE** and the file descriptor is already in the pollset, **pollset\_ctl** removes the file descriptor from the pollset. If the file descriptor is not already in the pollset, then **PS\_DELETE** causes **pollset\_ctl** to return an error.

The **pollset\_query** interface can be used to determine information about a file descriptor with respect to the pollset. If the file descriptor is in the pollset, **pollset\_query** returns 1 and *events* is set to the currently monitored events.

The **pollset\_poll** subroutine determines which file descriptors in the pollset that have events pending. The *polldata\_array* parameter contains a buffer address where **pollfd** structures are returned for file descriptors that have pending events. The number of events returned by a poll is limited by *array\_length*. The *timeout* parameter specifies the amount of time to wait if no events are pending. Setting *timeout* to 0 guarantees that the **pollset\_poll** subroutine returns immediately. Setting *timeout* to -1 specifies an infinite timeout. Other nonzero positive values specify the time to wait in milliseconds.

When events are returned from a **pollset\_poll** operation, each **pollfd** structure contains an *fd* member with the file descriptor set, an *events* member with the requested events, and an *revents* member with the events that have occurred.

A single pollset can be accessed by multiple threads in a multithreaded process. When multiple threads are polling one pollset and an event occurs for a file descriptor, only one thread can be prompted to receive the event. After a file descriptor is returned to a thread, new events will not be generated until the next **pollset\_poll** call. This behavior prevents all threads from being prompted on each event. Multiple threads can perform **pollset\_poll** operations at one time, but modifications to the pollset require exclusive access. A thread that tries to modify the pollset is blocked until all threads currently in poll operations have exited **pollset\_poll**. In addition, a thread calling **pollset\_destroy** is blocked until all threads have left the other system calls (**pollset\_ctl**, **pollset\_query**, and **pollset\_poll**).

A process can call **fork** after calling **pollset\_create**. The child process will already have a pollset ID per pollset, but **pollset\_destroy**, **pollset\_ctl**, **pollset\_query**, and **pollset\_poll** operations will fail with an **errno** value of **EACCES**.

After a file descriptor is added to a pollset, the file descriptor will not be removed until a **pollset\_ctl** call with the *cmd* of **PS\_DELETE** is executed. The file descriptor remains in the pollset even if the file descriptor is closed. A **pollset\_poll** operation on a pollset containing a closed file descriptor returns a **POLLNVAL** event for that file descriptor. If the file descriptor is later allocated to a new object, the new object will be polled on future **pollset\_poll** calls.

## Parameters

<i>array_length</i>	Specifies the length of the array parameters.
<i>maxfd</i>	Specifies the maximum number of file descriptors that can belong to the pollset.
<i>pollctl_array</i>	The pointer to an array of <b>poll_ctl</b> structures that describes the file descriptors (through the <b>pollfd</b> structure) and the unique operation to perform on each file descriptor (add, remove, or modify).
<i>polldata_array</i>	Returns the requested events that have occurred on the pollset.
<i>pollfd_query</i>	Points to a file descriptor that might or might not belong to the pollset. If it belongs to the pollset, then the requested events field of this parameter is updated to reflect what is currently being monitored for this file descriptor.
<i>ps</i>	Specifies the pollset ID.
<i>timeout</i>	Specifies the amount of time in milliseconds to wait for any monitored events to occur. A value of -1 blocks until some monitored event occurs.

## Return Values

Upon success, the **pollset\_destroy** returns 0. Upon failure, the **pollset\_destroy** subroutine returns -1 with **errno** set to the appropriate code.

Upon success, the **pollset\_create** subroutine returns a pollset ID of type **pollset\_t**. Upon failure, this system call returns -1 with **errno** set appropriately.

Upon success, **pollset\_ctl** returns 0. Upon failure, **pollset\_ctl** returns the 0-based problem element number of the **pollctl\_array** (for example, 2 is returned for element 3). If the first element is the problem element, or some other error occurs prior to processing the array of elements, -1 is returned and **errno** is set to the appropriate code. The calling application must acknowledge that elements in the array prior to the problem element were successfully processed and should attempt to call **pollset\_ctl** again with the elements of **pollctl\_array** beyond the problematic element.

If a file descriptor is not a member of the pollset, **pollset\_query** returns 0. If the file descriptor is in the pollset, **pollset\_query** returns 1 and *events* is set to the currently monitored events. If an error occurs after there is an attempt to determine if the file descriptor is a member of the pollset, then **pollset\_query** returns -1 with **errno** set to the appropriate return code.

The **pollset\_poll** subroutine returns the number of file descriptors on which requested events occurred. When no requested events occurred on any of the file descriptors, 0 is returned. A value of -1 is returned when an error occurs and **errno** is set to the appropriate code.

## Error Codes

<b>EACCES</b>	Process does not have permission to access a pollset.
<b>EAGAIN</b>	System resource temporarily not available.
<b>EFAULT</b>	Address supplied was not valid.
<b>EINTR</b>	A signal was received during the system call.
<b>EINVAL</b>	Invalid parameter.
<b>ENOMEM</b>	Insufficient system memory available.
<b>ENOSPC</b>	Maximum number of pollsets in use.
<b>EPERM</b>	Process does not have permission to create a pollset.

## Related Information

The “poll Subroutine” on page 1138.

---

## popen Subroutine

### Purpose

Initiates a pipe to a process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
FILE *popen ( Command, Type)  
const char *Command, *Type;
```

### Description

The **popen** subroutine creates a pipe between the calling program and a shell command to be executed.

**Note:** The **popen** subroutine runs only **sh** shell commands. The results are unpredictable if the *Command* parameter is not a valid **sh** shell command. If the terminal is in a trusted state, the **tsh** shell commands are run.

If streams opened by previous calls to the **popen** subroutine remain open in the parent process, the **popen** subroutine closes them in the child process.

The **popen** subroutine returns a pointer to a **FILE** structure for the stream.

**Attention:** If the original processes and the process started with the **popen** subroutine concurrently read or write a common file, neither should use buffered I/O. If they do, the results are unpredictable.

Some problems with an output filter can be prevented by flushing the buffer with the **fflush** subroutine.

### Parameters

*Command* Points to a null-terminated string containing a shell command line.  
*Type* Points to a null-terminated string containing an I/O mode. If the *Type* parameter is the value **r**, you can read from the standard output of the command by reading from the file *Stream*. If the *Type* parameter is the value **w**, you can write to the standard input of the command by writing to the file *Stream*.

Because open files are shared, a type **r** command can be used as an input filter and a type **w** command as an output filter.

### Return Values

The **popen** subroutine returns a null pointer if files or processes cannot be created, or if the shell cannot be accessed.

### Error Codes

The **popen** subroutine may set the **EINVAL** variable if the *Type* parameter is not valid. The **popen** subroutine may also set **errno** global variables as described by the **fork** or **pipe** subroutines.



## Related Information

The **fclose** or **fflush** (“fclose or fflush Subroutine” on page 253) subroutine, **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fork** or **vfork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **pclose** (“pclose Subroutine” on page 1002) subroutine, **pipe** (“pipe Subroutine” on page 1025) subroutine, **wait**, **waitpid**, or **wait3** subroutine.

Input and Output Handling.

File Systems and Directories in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## posix\_fadvise Subroutine

### Purpose

Provides advisory information to the system regarding future behavior of the application with respect to a given file.

### Syntax

```
#include <fcntl.h>
int posix_fadvise (int fd, off_t offset, size_t len, int advice);
```

### Description

This function advises the system on the expected future behavior of the application with regards to a given file. The system can take this advice into account when performing operations on file data specified by this function. The advice is given over the range covered by the *offset* parameter and continuing for the number of bytes specified by the *len* parameter. If the value of the *len* parameter is 0, then all data following the *offset* parameter is covered.

The *advice* parameter must be one of the following:

- **POSIX\_FADV\_NORMAL**
- **POSIX\_FADV\_SEQUENTIAL**
- **POSIX\_FADV\_RANDOM**
- **POSIX\_FADV\_WILLNEED**
- **POSIX\_FADV\_DONTNEED**
- **POSIX\_FADV\_NOREUSE**

### Parameters

<i>fd</i>	File descriptor of the file to be advised
<i>offset</i>	Represents the beginning of the address range
<i>len</i>	Determines the length of the address range
<i>advice</i>	Defines the advice to be given

### Return Values

Upon successful completion, the **posix\_fadvise** subroutine returns 0. Otherwise, one of the following error codes will be returned.

### Error Codes

<b>EBADF</b>	The <i>fd</i> parameter is not a valid file descriptor
--------------	--

EINVAL	The value of the <i>advice</i> parameter is invalid
ESPIPE	The <i>fd</i> parameter is associated with a pipe of FIFO

## Related Information

“*posix\_fallocate* Subroutine,” “*posix\_madvise* Subroutine” on page 1147, “*posix\_memalign*” on page 784.

---

## posix\_fallocate Subroutine

### Purpose

Reserve storage space for a given file descriptor.

### Syntax

```
#include <fcntl.h>
int posix_fallocate (int fd, off_t offset, size_t len);
```

### Description

This function reserves adequate space on the file system for the file data range beginning at the value specified by the *offset* parameter and continuing for the number of bytes specified by the *len* parameter. Upon successful return, subsequent writes to this file data range will not fail due to lack of free space on the file system media. Space allocated by the **posix\_fallocate** subroutine can be freed by a successful call to the **creat** subroutine or **open** subroutine, or by the **ftruncate** subroutine, which truncates the file size to a size smaller than the sum of the *offset* parameter and the *len* parameter.

### Parameters

<i>fd</i>	File descriptor of the file to reserve
<i>offset</i>	Represents the beginning of the address range
<i>len</i>	Determines the length of the address range

### Return Values

Upon successful completion, the **posix\_fallocate** subroutine returns 0. Otherwise, one of the following error codes will be returned.

### Error Codes

EBADF	The <i>fd</i> parameter is not a valid file descriptor
EBADF	The <i>fd</i> parameter references a file that was opened without write permission.
EFBIG	The value of the <i>offset</i> parameter plus the <i>len</i> parameter is greater than the maximum file size
EINTR	A signal was caught during execution
EIO	An I/O error occurred while reading from or writing to a file system
ENODEV	The <i>fd</i> parameter does not refer to a regular file.
EINVAL	The value of the <i>advice</i> parameter is invalid
ENOSPC	There is insufficient free space remaining on the file system storage media
ESPIPE	The <i>fd</i> parameter is associated with a pipe of FIFO

## Related Information

“`posix_fadvise` Subroutine” on page 1145, “`posix_madvise` Subroutine,” “`posix_memalign`” on page 784.

---

## posix\_madvise Subroutine

### Purpose

Provides advisory information to the system regarding future behavior of the application with respect to a given range of memory.

### Syntax

```
#include <sys/mman.h>
int posix_madvise (void *addr, size_t len, int advice);
```

### Description

This function advises the system on the expected future behavior of the application with regard to a given range of memory. The system can take this advice into account when performing operations on the data in memory specified by this function. The advice is given over the range covered by the *offset* parameter and continuing for the number of bytes specified by the *addr* parameter and continuing for the number of bytes specified by the *len* parameter.

The *advice* parameter must be one of the following:

- **POSIX\_MADV\_NORMAL**
- **POSIX\_MADV\_SEQUENTIAL**
- **POSIX\_MADV\_RANDOM**
- **POSIX\_MADV\_WILLNEED**
- **POSIX\_MADV\_DONTNEED**

### Parameters

*addr* Defines the beginning of the range of memory to be advised  
*len* Determines the length of the address range  
*advice* Defines the advice to be given

### Return Values

Upon successful completion, the **posix\_fadvise** subroutine returns 0. Otherwise, one of the following error codes will be returned.

### Error Codes

<b>EINVAL</b>	The value of the <i>advice</i> parameter is invalid
<b>ENOMEM</b>	Addresses in the range specified by the <i>addr</i> parameter and the <i>len</i> parameter are partially or completely outside the range of the process's address space.

## Related Information

“`posix_fallocate` Subroutine” on page 1146, “`posix_fadvise` Subroutine” on page 1145, “`posix_memalign`” on page 784.

---

## posix\_openpt Subroutine

### Purpose

Opens a pseudo-terminal device.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <stdlib.h>
#include <fcntl.h>

int posix_openpt (oflag)
)
int oflag;
```

### Description

The **posix\_openpt** subroutine establishes a connection between a master device for a pseudo terminal and a file descriptor. The file descriptor is used by other I/O functions that refer to that pseudo terminal.

The file status flags and file access modes of the open file description are set according to the value of the *oflag* parameter.

### Parameters

*oflag*

Values for the *oflag* parameter are constructed by a bitwise-inclusive OR of flags from the following list, defined in the **<fcntl.h>** file:

#### **O\_RDWR**

Open for reading and writing.

#### **O\_NOCTTY**

If set, the **posix\_openpt** subroutine does not cause the terminal device to become the controlling terminal for the process.

The behavior of other values for the *oflag* parameter is unspecified.

### Return Values

Upon successful completion, the **posix\_openpt** subroutine opens a master pseudo-terminal device and returns a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **posix\_openpt** subroutine will fail if:

**EMFILE**

OPEN\_MAX file descriptors are currently open in the calling process.

**ENFILE**

The maximum allowable number of files is currently open in the system.

The `posix_openpt` subroutine may fail if:

<b>EINVAL</b>	The value of the <i>oflag</i> parameter is not valid.
<b>EAGAIN</b>	Out of pseudo-terminal resources.
<b>ENOSR</b>	Out of STREAMS resources.

## Examples

The following example describes how to open a pseudo-terminal and return the name of the slave device and file descriptor

```
#include <fcntl.h>
#include <stdio.h>

int masterfd, slavefd;
char *slavedevice;

masterfd = posix_openpt(O_RDWR|O_NOCTTY);

if (masterfd == -1
    || grantpt (masterfd) == -1
    || unlockpt (masterfd) == -1
    || (slavedevice = ptsname (masterfd)) == NULL)
    return -1;

printf("slave device is: %s\n", slavedevice);

slavefd = open(slave, O_RDWR|O_NOCTTY);
if (slavefd < 0)
    return -1;
```

## Related Information

“grantpt Subroutine” on page 485, “open, openx, open64, creat, or creat64 Subroutine” on page 935, “ptsname Subroutine” on page 1320.

**unlockpt** Subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

`<fcntl.h>` file in *AIX 5L Version 5.3 Files Reference*.

---

## posix\_spawn or posix\_spawnp Subroutine

### Purpose

Spawns a process.

### Syntax

```
int posix_spawn(pid_t *restrict pid, const char *restrict path,
               const posix_spawn_file_actions_t *file_actions,
               const posix_spawnattr_t *restrict attrp,
               char *const argv[restrict], char *const envp[restrict]);
int posix_spawnp(pid_t *restrict pid, const char *restrict file,
                const posix_spawn_file_actions_t *file_actions,
                const posix_spawnattr_t *restrict attrp,
                char *const argv[restrict], char * const envp[restrict]);
```

## Description

The **posix\_spawn** and **posix\_spawnnp** subroutines create a new process (child process) from the specified process image. The new process image is constructed from a regular executable file called the *new process image file*.

When a C program is executed as the result of this call, the program is entered as a C-language function call as follows:

```
int main(int argc, char *argv[]);
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings.

The *argv* parameter is an array of character pointers to null-terminated strings. The last member of this array is a null pointer and is not counted in *argc*. These strings constitute the argument list available to the new process image. The value in *argv*[0] should point to a file name that is associated with the process image being started by the **posix\_spawn** or **posix\_spawnnp** function.

The argument *envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new process image. The environment array is terminated by a null pointer.

The number of bytes available for the child process' combined argument and environment lists is **{ARG\_MAX}**. The implementation specifies in the system documentation whether any list overhead, such as length words, null terminators, pointers, or alignment bytes, is included in this total.

The path argument to **posix\_spawn** is a path name that identifies the new process image file to execute.

The file parameter to **posix\_spawnnp** is used to construct a path name that identifies the new process image file. If the file parameter contains a slash character (/), the file parameter is used as the path name for the new process image file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable **PATH**. If this environment variable is not defined, the results of the search are implementation-defined.

If *file\_actions* is a null pointer, file descriptors that are open in the calling process remain open in the child process, except for those whose **FD\_CLOEXEC** flag is set (see “fcntl, dup, or dup2 Subroutine” on page 255). For those file descriptors that remain open, all attributes of the corresponding open file descriptions, including file locks, remain unchanged.

If *file\_actions* is not a null pointer, the file descriptors open in the child process are those open in the calling process as modified by the spawn file actions object pointed to by *file\_actions* and the **FD\_CLOEXEC** flag of each remaining open file descriptor after the spawn file actions have been processed. The effective order of processing the spawn file actions is as follows:

1. The set of open file descriptors for the child process is initially the same set as is open for the calling process. All attributes of the corresponding open file descriptions, including file locks (see “fcntl, dup, or dup2 Subroutine” on page 255), remain unchanged.
2. The signal mask, signal default actions, and the effective user and group IDs for the child process are changed as specified in the attributes object referenced by *attrp*.
3. The file actions specified by the spawn file actions object are performed in the order in which they were added to the spawn file actions object.
4. Any file descriptor that has its **FD\_CLOEXEC** flag set is closed.

The **posix\_spawnattr\_t** spawn attributes object type is defined in the **spawn.h** header file. Its attributes are defined as follows:

- If the **POSIX\_SPAWN\_SETPGROUP** flag is set in the **spawn-flags** attribute of the object referenced by *attrp*, and the **spawn-pgroup** attribute of the same object is non-zero, the child's process group is as specified in the **spawn-pgroup** attribute of the object referenced by *attrp*.
- As a special case, if the **POSIX\_SPAWN\_SETPGROUP** flag is set in the **spawn-flags** attribute of the object referenced by *attrp*, and the **spawn-pgroup** attribute of the same object is set to 0, then the child is in a new process group with a process group ID equal to its process ID.
- If the **POSIX\_SPAWN\_SETPGROUP** flag is not set in the **spawn-flags** attribute of the object referenced by *attrp*, the new child process inherits the parent's process group.
- If the **POSIX\_SPAWN\_SETSCHEDPARAM** flag is set in the **spawn-flags** attribute of the object referenced by *attrp*, but **POSIX\_SPAWN\_SETSCHEDULER** is not set, the new process image initially has the scheduling policy of the calling process with the scheduling parameters specified in the **spawn-schedparam** attribute of the object referenced by *attrp*.
- If the **POSIX\_SPAWN\_SETSCHEDULER** flag is set in the **spawn-flags** attribute of the object referenced by *attrp* (regardless of the setting of the **POSIX\_SPAWN\_SETSCHEDPARAM** flag), the new process image initially has the scheduling policy specified in the **spawn-schedpolicy** attribute of the object referenced by *attrp* and the scheduling parameters specified in the **spawn-schedparam** attribute of the same object.
- The **POSIX\_SPAWN\_RESETIDS** flag in the **spawn-flags** attribute of the object referenced by *attrp* governs the effective user ID of the child process. If this flag is not set, the child process inherits the parent process' effective user ID. If this flag is set, the child process' effective user ID is reset to the parent's real user ID. In either case, if the set-user-ID mode bit of the new process image file is set, the effective user ID of the child process becomes that file's owner ID before the new process image begins execution.
- The **POSIX\_SPAWN\_RESETIDS** flag in the **spawn-flags** attribute of the object referenced by *attrp* also governs the effective group ID of the child process. If this flag is not set, the child process inherits the parent process' effective group ID. If this flag is set, the child process' effective group ID is reset to the parent's real group ID. In either case, if the set-group-ID mode bit of the new process image file is set, the effective group ID of the child process becomes that file's group ID before the new process image begins execution.
- If the **POSIX\_SPAWN\_SETSIGMASK** flag is set in the **spawn-flags** attribute of the object referenced by *attrp*, the child process initially has the signal mask specified in the **spawn-sigmask** attribute of the object referenced by *attrp*.
- If the **POSIX\_SPAWN\_SETSIGDEF** flag is set in the **spawn-flags** attribute of the object referenced by *attrp*, the signals specified in the **spawn-sigdefault** attribute of the same object is set to their default actions in the child process. Signals set to the default action in the parent process are set to the default action in the child process. Signals set to be caught by the calling process are set to the default action in the child process.
- Except for **SIGCHLD**, signals set to be ignored by the calling process image are set to be ignored by the child process, unless otherwise specified by the **POSIX\_SPAWN\_SETSIGDEF** flag being set in the **spawn-flags** attribute of the object referenced by *attrp* and the signals being indicated in the **spawn-sigdefault** attribute of the object referenced by *attrp*.
- If the **SIGCHLD** signal is set to be ignored by the calling process, it is unspecified whether the **SIGCHLD** signal is set to be ignored or set to the default action in the child process. This is true unless otherwise specified by the **POSIX\_SPAWN\_SETSIGDEF** flag being set in the **spawn-flags** attribute of the object referenced by *attrp* and the **SIGCHLD** signal being indicated in the **spawn-sigdefault** attribute of the object referenced by *attrp*.
- If the value of the *attrp* pointer is NULL, then the default values are used.

All process attributes, other than those influenced by the attributes set in the object referenced by *attrp* in the preceding list or by the file descriptor manipulations specified in *file\_actions*, are displayed in the new process image as though **fork** had been called to create a child process and then a member of the **exec** family of functions had been called by the child process to execute the new process image.

By default, fork handlers are not run in **posix\_spawn** or **posix\_spawnnp** routines. To enable fork handlers, set the **POSIX\_SPAWN\_FORK\_HANDLERS** flag in the attribute.

## Return Values

Upon successful completion, **posix\_spawn** and **posix\_spawnnp** return the process ID of the child process to the parent process, in the variable pointed to by a non-NULL *pid* argument, and return 0 as the function return value. Otherwise, no child process is created, the value stored into the variable pointed to by a non-NULL *pid* is unspecified, and an error number is returned as the function return value to indicate the error. If the *pid* argument is a null pointer, the process ID of the child is not returned to the caller.

## Error Codes

The **posix\_spawn** and **posix\_spawnnp** subroutines will fail if the following is true:

**EINVAL**            The value specified by *file\_actions* or *attrp* is invalid.

The error codes for the **posix\_spawn** and **posix\_spawnnp** subroutines are affected by the following conditions:

- If this error occurs after the calling process successfully returns from the **posix\_spawn** or **posix\_spawnnp** function, the child process might exit with exit status 127.
- If **posix\_spawn** or **posix\_spawnnp** fail for any of the reasons that would cause **fork** or one of the **exec** family of functions to fail, an error value is returned as described by **fork** and **exec**, respectively (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).
- If **POSIX\_SPAWN\_SETPGROUP** is set in the **spawn-flags** attribute of the object referenced by *attrp*, and **posix\_spawn** or **posix\_spawnnp** fails while changing the child's process group, an error value is returned as described by **setpgid** (or, if the error occurs after the calling process successfully returns, the child process shall exit with exit status 127).
- If **POSIX\_SPAWN\_SETSCHEDPARAM** is set and **POSIX\_SPAWN\_SETSCHEDULER** is not set in the **spawn-flags** attribute of the object referenced by *attrp*, then if **posix\_spawn** or **posix\_spawnnp** fails for any of the reasons that would cause **sched\_setparam** to fail, an error value is returned as described by **sched\_setparam** (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).
- If **POSIX\_SPAWN\_SETSCHEDULER** is set in the **spawn-flags** attribute of the object referenced by *attrp*, and if **posix\_spawn** or **posix\_spawnnp** fails for any of the reasons that would cause **sched\_setscheduler** to fail, an error value is returned as described by **sched\_setscheduler** (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).
- If the *file\_actions* argument is not NULL and specifies any **close**, **dup2**, or **open** actions to be performed, and if **posix\_spawn** or **posix\_spawnnp** fails for any of the reasons that would cause **close**, **dup2**, or **open** to fail, an error value is returned as described by **close**, **dup2**, and **open**, respectively (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127). An open file action might, by itself, result in any of the errors described by **close** or **dup2**, in addition to those described by **open**.

## Related Information

The “getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine” on page 386, “chmod or fchmod Subroutine” on page 148, “close Subroutine” on page 175, “fcntl, dup, or dup2 Subroutine” on page 255, “exec: execl, execl, execlp, execv, execve, execvp, or exect Subroutine” on page 236, “exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243, “fork, f\_fork, or vfork Subroutine” on page 291, “kill or killpg Subroutine” on page 580, “open, openx, open64, creat, or creat64 Subroutine” on page 935, “posix\_spawn\_file\_actions\_addclose or posix\_spawn\_file\_actions\_addopen Subroutine” on page 1153, “posix\_spawn\_file\_actions\_adddup2 Subroutine” on page 1154, “posix\_spawn\_file\_actions\_destroy or posix\_spawn\_file\_actions\_init Subroutine” on page 1155,



“`posix_spawnattr_destroy` or `posix_spawnattr_init` Subroutine” on page 1156,  
“`posix_spawnattr_getsigdefault` or `posix_spawnattr_setsigdefault` Subroutine” on page 1161,  
“`posix_spawnattr_getflags` or `posix_spawnattr_setflags` Subroutine” on page 1157,  
“`posix_spawnattr_getpgroup` or `posix_spawnattr_setpgroup` Subroutine” on page 1158,  
“`posix_spawnattr_getschedparam` or `posix_spawnattr_setschedparam` Subroutine” on page 1159,  
“`posix_spawnattr_getschedpolicy` or `posix_spawnattr_setschedpolicy` Subroutine” on page 1160,  
“`posix_spawnattr_getsigmask` or `posix_spawnattr_setsigmask` Subroutine” on page 1162.

---

## `posix_spawn_file_actions_addclose` or `posix_spawn_file_actions_addopen` Subroutine

### Purpose

Adds **close** or **open** action to the spawn file actions object.

### Syntax

```
#include <spawn.h>

int posix_spawn_file_actions_addclose(posix_spawn_file_actions_t *
    file_actions, int fildes);
int posix_spawn_file_actions_addopen(posix_spawn_file_actions_t *
    restrict file_actions, int fildes,
    const char *restrict path, int oflag, mode_t mode);
```

### Description

The `posix_spawn_file_actions_addclose` and `posix_spawn_file_actions_addopen` subroutines **close** or **open** action to a spawn file actions object.

A spawn file actions object is of type `posix_spawn_file_actions_t` (defined in the `spawn.h` header file) and is used to specify a series of actions to be performed by a `posix_spawn` or `posix_spawnnp` operation in order to arrive at the set of **open** file descriptors for the child process given the set of open file descriptors of the parent. Comparison or assignment operators for the type `posix_spawn_file_actions_t` are not defined.

A spawn file actions object, when passed to `posix_spawn` or `posix_spawnnp`, specifies how the set of **open** file descriptors in the calling process is transformed into a set of potentially **open** file descriptors for the spawned process. This transformation is as if the specified sequence of actions was performed exactly once, in the context of the spawned process (prior to running the new process image), in the order in which the actions were added to the object. Additionally, when the new process image is run, any file descriptor (from this new set) that has its **FD\_CLOEXEC** flag set is closed (see “`posix_spawn` or `posix_spawnnp` Subroutine” on page 1149).

The `posix_spawn_file_actions_addclose` function adds a **close** action to the object referenced by `file_actions` that causes the file descriptor `fildes` to be closed (as if `close(fildes)` had been called) when a new process is spawned using this file actions object.

The `posix_spawn_file_actions_addopen` function adds an **open** action to the object referenced by `file_actions` that causes the file named by `path` to be opened, as if `open(path, oflag, mode)` had been called, and the returned file descriptor, if not `fildes`, had been changed to `fildes` when a new process is spawned using this file actions object. If `fildes` was already an **open** file descriptor, it closes before the new file is opened.

The string described by `path` is copied by the `posix_spawn_file_actions_addopen` function.

## Return Values

Upon successful completion, the **posix\_spawn\_file\_actions\_addclose** and **posix\_spawn\_file\_actions\_addopen** subroutines return 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawn\_file\_actions\_addclose** and **posix\_spawn\_file\_actions\_addopen** subroutines fail if the following is true:

**EBADF**                The value specified by *fildev* is negative, or greater than or equal to **{OPEN\_MAX}**.

The **posix\_spawn\_file\_actions\_addclose** and **posix\_spawn\_file\_actions\_addopen** subroutines might fail if the following are true:

**EINVAL**              The value specified by *file\_actions* is invalid.

**ENOMEM**             Insufficient memory exists to add to the spawn file actions object.

It is not an error for the *fildev* argument passed to these functions to specify a file descriptor for which the specified operation could not be performed at the time of the call. Any such error will be detected when the associated file actions object is used later during a **posix\_spawn** or **posix\_spawnnp** operation.

## Related Information

The “close Subroutine” on page 175, “fcntl, dup, or dup2 Subroutine” on page 255, “open, openx, open64, creat, or creat64 Subroutine” on page 935, “posix\_spawn or posix\_spawnnp Subroutine” on page 1149, “posix\_spawn\_file\_actions\_adddup2 Subroutine,” “posix\_spawn\_file\_actions\_destroy or posix\_spawn\_file\_actions\_init Subroutine” on page 1155.

---

## posix\_spawn\_file\_actions\_adddup2 Subroutine

### Purpose

Adds **dup2** action to the spawn file actions object.

### Syntax

```
#include <spawn.h>
```

```
int posix_spawn_file_actions_adddup2(posix_spawn_file_actions_t *  
    file_actions, int fildev, int newfildev);
```

### Description

The **posix\_spawn\_file\_actions\_adddup2** subroutine adds a **dup2** action to the object referenced by *file\_actions* that causes the file descriptor *fildev* to be duplicated as *newfildev* when a new process is spawned using this file actions object. This functions as if **dup2( fildev, newfildev)** had been called.

A spawn file actions object is as defined in **posix\_spawn\_file\_actions\_addclose**.

## Return Values

Upon successful completion, the **posix\_spawn\_file\_actions\_adddup2** subroutine returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawn\_file\_actions\_adddup2** subroutine will fail if the following are true:

<b>EBADF</b>	The value specified by <i>fildes</i> or <i>newfildes</i> is negative, or greater than or equal to <b>{OPEN_MAX}</b> .
<b>ENOMEM</b>	Insufficient memory exists to add to the spawn file actions object.

The **posix\_spawn\_file\_actions\_adddup2** subroutine might fail if the following is true:

<b>EINVAL</b>	The value specified by <i>file_actions</i> is invalid.
---------------	--

It is not an error for the *fildes* argument passed to this subroutine to specify a file descriptor for which the specified operation could not be performed at the time of the call. Any such error will be detected when the associated file actions object is used later during a **posix\_spawn** or **posix\_spawnp** operation.

## Related Information

The “fcntl, dup, or dup2 Subroutine” on page 255, “posix\_spawn or posix\_spawnp Subroutine” on page 1149, “posix\_spawn\_file\_actions\_addclose or posix\_spawn\_file\_actions\_addopen Subroutine” on page 1153, “posix\_spawn\_file\_actions\_destroy or posix\_spawn\_file\_actions\_init Subroutine.”

---

## posix\_spawn\_file\_actions\_destroy or posix\_spawn\_file\_actions\_init Subroutine

### Purpose

Destroys and initializes a spawn file actions object.

### Syntax

```
#include <spawn.h>

int posix_spawn_file_actions_destroy(posix_spawn_file_actions_t *
    file_actions);
int posix_spawn_file_actions_init(posix_spawn_file_actions_t *
    file_actions);
```

### Description

The **posix\_spawn\_file\_actions\_destroy** subroutine destroys the object referenced by *file\_actions*; the object becomes, in effect, uninitialized. An implementation can cause **posix\_spawn\_file\_actions\_destroy** to set the object referenced by *file\_actions* to an invalid value. A destroyed spawn file actions object can be reinitialized using **posix\_spawn\_file\_actions\_init**; the results of otherwise referencing the object after it has been destroyed are undefined.

The **posix\_spawn\_file\_actions\_init** function initializes the object referenced by *file\_actions* to contain no file actions for **posix\_spawn** or **posix\_spawnp** to perform.

A spawn file actions object is as defined in **posix\_spawn\_file\_actions\_addclose**. The effect of initializing a previously initialized spawn file actions object is undefined.

### Return Values

Upon successful completion, the **posix\_spawn\_file\_actions\_destroy** and **posix\_spawn\_file\_actions\_init** subroutines return 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawn\_file\_actions\_init** subroutine will fail if the following is true:

**ENOMEM**            Insufficient memory exists to initialize the spawn file actions object.

The **posix\_spawn\_file\_actions\_destroy** subroutine might fail if the following is true:

**EINVAL**            The value specified by *file\_actions* is invalid.

## Related Information

The “[posix\\_spawn or posix\\_spawn Subroutine](#)” on page 1149.

---

## posix\_spawnattr\_destroy or posix\_spawnattr\_init Subroutine

### Purpose

Destroys and initializes a spawn attributes object.

### Syntax

```
#include <spawn.h>
```

```
int posix_spawnattr_destroy(posix_spawnattr_t *attr);  
int posix_spawnattr_init(posix_spawnattr_t *attr);
```

### Description

The **posix\_spawnattr\_destroy** subroutine destroys a spawn attributes object. A destroyed *attr* attributes object can be reinitialized using **posix\_spawnattr\_init**; the results of otherwise referencing the object after it has been destroyed are undefined. An implementation can cause **posix\_spawnattr\_destroy** to set the object referenced by *attr* to an invalid value.

The **posix\_spawnattr\_init** subroutine initializes a spawn attributes object *attr* with the default value for all of the individual attributes used by the implementation. Results are undefined if **posix\_spawnattr\_init** is called specifying an *attr* attributes object that is already initialized.

A spawn attributes object is of type **posix\_spawnattr\_t** (defined in the **spawn.h** header file) and is used to specify the inheritance of process attributes across a spawn operation. Comparison or assignment operators for the type **posix\_spawnattr\_t** are not defined.

Each implementation documents the individual attributes it uses and their default values unless these values are defined by IEEE Std 1003.1-2001. Attributes not defined by IEEE Std 1003.1-2001, their default values, and the names of the associated functions to get and set those attribute values are implementation-defined.

The resulting spawn attributes object (possibly modified by setting individual attribute values), is used to modify the behavior of **posix\_spawn** or **posix\_spawnnp**. After a spawn attributes object has been used to spawn a process by a call to a **posix\_spawn** or **posix\_spawnnp**, any function affecting the attributes object (including destruction) will not affect any process that has been spawned in this way.

### Return Values

Upon successful completion, the **posix\_spawnattr\_destroy** and **posix\_spawnattr\_init** subroutines return 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The `posix_spawnattr_destroy` subroutine might fail if the following is true:

**EINVAL**            The value specified by *attr* is invalid.

## Related Information

The “`posix_spawn` or `posix_spawnp` Subroutine” on page 1149, “`posix_spawnattr_getsigdefault` or `posix_spawnattr_setsigdefault` Subroutine” on page 1161, “`posix_spawnattr_getflags` or `posix_spawnattr_setflags` Subroutine,” “`posix_spawnattr_getpgroup` or `posix_spawnattr_setpgroup` Subroutine” on page 1158, “`posix_spawnattr_getschedparam` or `posix_spawnattr_setschedparam` Subroutine” on page 1159, “`posix_spawnattr_getschedpolicy` or `posix_spawnattr_setschedpolicy` Subroutine” on page 1160, “`posix_spawnattr_getsigmask` or `posix_spawnattr_setsigmask` Subroutine” on page 1162.

---

## `posix_spawnattr_getflags` or `posix_spawnattr_setflags` Subroutine

### Purpose

Gets and sets the **spawn-flags** attribute of a spawn attributes object.

### Syntax

```
#include <spawn.h>

int posix_spawnattr_getflags(const posix_spawnattr_t *restrict attr,
                             short *restrict flags);
int posix_spawnattr_setflags(posix_spawnattr_t *attr, short flags);
```

### Description

The `posix_spawnattr_getflags` subroutine obtains the value of the **spawn-flags** attribute from the attributes object referenced by *attr*. The `posix_spawnattr_setflags` subroutine sets the **spawn-flags** attribute in an initialized attributes object referenced by *attr*. The **spawn-flags** attribute is used to indicate which process attributes are to be changed in the new process image when invoking `posix_spawn` or `posix_spawnp`. It is the bitwise-inclusive OR of 0 or more of the following flags:

- **POSIX\_SPAWN\_RESETIDS**
- **POSIX\_SPAWN\_SETPGROUP**
- **POSIX\_SPAWN\_SETSIGDEF**
- **POSIX\_SPAWN\_SETSIGMASK**
- **POSIX\_SPAWN\_SETSCHEDPARAM**
- **POSIX\_SPAWN\_SETSCHEDULER**

These flags are defined in the `spawn.h` header file. The default value of this attribute is as if no flags were set.

### Return Values

Upon successful completion, the `posix_spawnattr_getflags` subroutine returns 0 and stores the value of the **spawn-flags** attribute of *attr* into the object referenced by the *flags* parameter; otherwise, an error number is returned to indicate the error.

Upon successful completion, the `posix_spawnattr_setflags` subroutine returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawnattr\_getflags** and **posix\_spawnattr\_setflags** subroutines will fail if the following is true:

**EINVAL**            The value specified by *attr* is invalid.

The **posix\_spawnattr\_setflags** subroutine might fail if the following is true:

**EINVAL**            The value of the attribute being set is not valid.

## Related Information

The “**posix\_spawn** or **posix\_spawnnp** Subroutine” on page 1149, “**posix\_spawn\_file\_actions\_addclose** or **posix\_spawn\_file\_actions\_addopen** Subroutine” on page 1153, “**posix\_spawn\_file\_actions\_adddup2** Subroutine” on page 1154, “**posix\_spawn\_file\_actions\_destroy** or **posix\_spawn\_file\_actions\_init** Subroutine” on page 1155, “**posix\_spawnattr\_destroy** or **posix\_spawnattr\_init** Subroutine” on page 1156, “**posix\_spawnattr\_getsigdefault** or **posix\_spawnattr\_setsigdefault** Subroutine” on page 1161, “**posix\_spawnattr\_getpgroup** or **posix\_spawnattr\_setpgroup** Subroutine,” “**posix\_spawnattr\_getschedparam** or **posix\_spawnattr\_setschedparam** Subroutine” on page 1159, “**posix\_spawnattr\_getschedpolicy** or **posix\_spawnattr\_setschedpolicy** Subroutine” on page 1160, “**posix\_spawnattr\_getsigmask** or **posix\_spawnattr\_setsigmask** Subroutine” on page 1162

---

## posix\_spawnattr\_getpgroup or posix\_spawnattr\_setpgroup Subroutine

### Purpose

Gets and sets the **spawn-pgroup** attribute of a spawn attributes object.

### Syntax

```
#include <spawn.h>
```

```
int posix_spawnattr_getpgroup(const posix_spawnattr_t *restrict attr,  
                             pid_t *restrict pgroup);  
int posix_spawnattr_setpgroup(posix_spawnattr_t *attr, pid_t pgroup);
```

### Description

The **posix\_spawnattr\_getpgroup** subroutine gets the value of the **spawn-pgroup** attribute from the attributes object referenced by *attr*.

The **posix\_spawnattr\_setpgroup** subroutine sets the **spawn-pgroup** attribute in an initialized attributes object referenced by *attr*.

The **spawn-pgroup** attribute represents the process group to be joined by the new process image in a spawn operation (if **POSIX\_SPAWN\_SETPGROUP** is set in the **spawn-flags** attribute). The default value of this attribute is 0.

### Return Values

Upon successful completion, the **posix\_spawnattr\_getpgroup** subroutine returns 0 and stores the value of the **spawn-pgroup** attribute of *attr* into the object referenced by the *pgroup* parameter; otherwise, an error number is returned to indicate the error.

Upon successful completion, the **posix\_spawnattr\_setpgroup** subroutine returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The `posix_spawnattr_getpgroup` and `posix_spawnattr_setpgroup` subroutines might fail if the following is true:

**EINVAL** The value specified by *attr* is invalid.

The `posix_spawnattr_setpgroup` subroutine might fail if the following is true:

**EINVAL** The value of the attribute being set is not valid.

## Related Information

The “`posix_spawn` or `posix_spawnnp` Subroutine” on page 1149, “`posix_spawn_file_actions_addclose` or `posix_spawn_file_actions_addopen` Subroutine” on page 1153, “`posix_spawn_file_actions_adddup2` Subroutine” on page 1154, “`posix_spawn_file_actions_destroy` or `posix_spawn_file_actions_init` Subroutine” on page 1155, “`posix_spawnattr_destroy` or `posix_spawnattr_init` Subroutine” on page 1156, “`posix_spawnattr_getsigdefault` or `posix_spawnattr_setsigdefault` Subroutine” on page 1161, “`posix_spawnattr_getflags` or `posix_spawnattr_setflags` Subroutine” on page 1157, “`posix_spawnattr_getschedparam` or `posix_spawnattr_setschedparam` Subroutine,” “`posix_spawnattr_getschedpolicy` or `posix_spawnattr_setschedpolicy` Subroutine” on page 1160, “`posix_spawnattr_getsigmask` or `posix_spawnattr_setsigmask` Subroutine” on page 1162

---

## `posix_spawnattr_getschedparam` or `posix_spawnattr_setschedparam` Subroutine

### Purpose

Gets and sets the **spawn-schedparam** attribute of a spawn attributes object.

### Syntax

```
#include <spawn.h>
#include <sched.h>

int posix_spawnattr_getschedparam(const posix_spawnattr_t *
    restrict attr, struct sched_param *restrict schedparam);
int posix_spawnattr_setschedparam(posix_spawnattr_t *restrict attr,
    const struct sched_param *restrict schedparam);
```

### Description

The `posix_spawnattr_getschedparam` subroutine gets the value of the **spawn-schedparam** attribute from the attributes object referenced by *attr*.

The `posix_spawnattr_setschedparam` subroutine sets the **spawn-schedparam** attribute in an initialized attributes object referenced by *attr*.

The **spawn-schedparam** attribute represents the scheduling parameters to be assigned to the new process image in a spawn operation (if **POSIX\_SPAWN\_SETSCHEDULER** or **POSIX\_SPAWN\_SETSCHEDPARAM** is set in the **spawn-flags** attribute). The default value of this attribute is unspecified.

### Return Values

Upon successful completion, the `posix_spawnattr_getschedparam` subroutine returns 0 and stores the value of the **spawn-schedparam** attribute of *attr* into the object referenced by the *schedparam* parameter; otherwise, an error number is returned to indicate the error.

Upon successful completion, the **posix\_spawnattr\_setschedparam** subroutine returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawnattr\_getschedparam** and **posix\_spawnattr\_setschedparam** subroutines might fail if the following is true:

**EINVAL**            The value specified by *attr* is invalid.

The **posix\_spawnattr\_setschedparam** subroutine might fail if the following is true:

**EINVAL**            The value of the attribute being set is not valid.

## Related Information

The “**posix\_spawn** or **posix\_spawnnp** Subroutine” on page 1149, “**posix\_spawn\_file\_actions\_addclose** or **posix\_spawn\_file\_actions\_addopen** Subroutine” on page 1153, “**posix\_spawn\_file\_actions\_adddup2** Subroutine” on page 1154, “**posix\_spawn\_file\_actions\_destroy** or **posix\_spawn\_file\_actions\_init** Subroutine” on page 1155, “**posix\_spawnattr\_destroy** or **posix\_spawnattr\_init** Subroutine” on page 1156, “**posix\_spawnattr\_getsigdefault** or **posix\_spawnattr\_setsigdefault** Subroutine” on page 1161, “**posix\_spawnattr\_getflags** or **posix\_spawnattr\_setflags** Subroutine” on page 1157, “**posix\_spawnattr\_getpgroup** or **posix\_spawnattr\_setpgroup** Subroutine” on page 1158, “**posix\_spawnattr\_getschedpolicy** or **posix\_spawnattr\_setschedpolicy** Subroutine,” “**posix\_spawnattr\_getsigmask** or **posix\_spawnattr\_setsigmask** Subroutine” on page 1162

---

## posix\_spawnattr\_getschedpolicy or posix\_spawnattr\_setschedpolicy Subroutine

### Purpose

Gets and sets the **spawn-schedpolicy** attribute of a spawn attributes object.

### Syntax

```
#include <spawn.h>
#include <sched.h>

int posix_spawnattr_getschedpolicy(const posix_spawnattr_t *
    restrict attr, int *restrict schedpolicy);
int posix_spawnattr_setschedpolicy(posix_spawnattr_t *attr,
    int schedpolicy);
```

### Description

The **posix\_spawnattr\_getschedpolicy** subroutine gets the value of the **spawn-schedpolicy** attribute from the attributes object referenced by *attr*.

The **posix\_spawnattr\_setschedpolicy** subroutine sets the **spawn-schedpolicy** attribute in an initialized attributes object referenced by *attr*.

The **spawn-schedpolicy** attribute represents the scheduling policy to be assigned to the new process image in a spawn operation (if **POSIX\_SPAWN\_SETSCHEDULER** is set in the **spawn-flags** attribute). The default value of this attribute is unspecified.



## Return Values

Upon successful completion, the **posix\_spawnattr\_getschedpolicy** subroutine returns 0 and stores the value of the **spawn-schedpolicy** attribute of *attr* into the object referenced by the *schedpolicy* parameter; otherwise, an error number is returned to indicate the error.

Upon successful completion, **posix\_spawnattr\_setschedpolicy** returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The following **posix\_spawnattr\_getschedpolicy** and **posix\_spawnattr\_setschedpolicy** subroutines might fail if the following is true:

**EINVAL**            The value specified by *attr* is invalid.

The **posix\_spawnattr\_setschedpolicy** subroutine might fail if the following is true:

**EINVAL**            The value of the attribute being set is not valid.

## Related Information

The “**posix\_spawn** or **posix\_spawnp** Subroutine” on page 1149, “**posix\_spawn\_file\_actions\_addclose** or **posix\_spawn\_file\_actions\_addopen** Subroutine” on page 1153, “**posix\_spawn\_file\_actions\_adddup2** Subroutine” on page 1154, “**posix\_spawn\_file\_actions\_destroy** or **posix\_spawn\_file\_actions\_init** Subroutine” on page 1155, “**posix\_spawnattr\_destroy** or **posix\_spawnattr\_init** Subroutine” on page 1156, “**posix\_spawnattr\_getsigdefault** or **posix\_spawnattr\_setsigdefault** Subroutine,” “**posix\_spawnattr\_getflags** or **posix\_spawnattr\_setflags** Subroutine” on page 1157, “**posix\_spawnattr\_getpgroup** or **posix\_spawnattr\_setpgroup** Subroutine” on page 1158, “**posix\_spawnattr\_getschedparam** or **posix\_spawnattr\_setschedparam** Subroutine” on page 1159, “**posix\_spawnattr\_getsigmask** or **posix\_spawnattr\_setsigmask** Subroutine” on page 1162.

---

## posix\_spawnattr\_getsigdefault or posix\_spawnattr\_setsigdefault Subroutine

### Purpose

Gets and sets the **spawn-sigdefault** attribute of a spawn attributes object.

### Syntax

```
#include <signal.h>
#include <spawn.h>

int posix_spawnattr_getsigdefault(const posix_spawnattr_t *
    restrict attr, sigset_t *restrict sigdefault);
int posix_spawnattr_setsigdefault(posix_spawnattr_t *restrict attr,
    const sigset_t *restrict sigdefault);
```

### Description

The **posix\_spawnattr\_getsigdefault** subroutine gets the value of the **spawn-sigdefault** attribute from the attributes object referenced by *attr*.

The **posix\_spawnattr\_setsigdefault** subroutine sets the **spawn-pgroup** attribute in an initialized attributes object referenced by *attr*.

The **spawn-sigdefault** attribute represents the set of signals to be forced to default signal handling in the new process image by a spawn operation (if **POSIX\_SPAWN\_SETSIGDEF** is set in the **spawn-flags** attribute). The default value of this attribute is an empty signal set.

## Return Values

Upon successful completion, the **posix\_spawnattr\_getsigdefault** subroutine returns 0 and stores the value of the **spawn-sigdefault** attribute of *attr* into the object referenced by the *sigdefault* parameter; otherwise, an error number is returned to indicate the error.

Upon successful completion, the **posix\_spawnattr\_setsigdefault** subroutine returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawnattr\_getsigdefault** and **posix\_spawnattr\_setsigdefault** subroutines might fail if the following is true:

**EINVAL**            The value specified by *attr* is invalid.

The **posix\_spawnattr\_setsigdefault** subroutine might fail if the following is true:

**EINVAL**            The value of the attribute being set is not valid.

## Related Information

The “**posix\_spawn** or **posix\_spawnp** Subroutine” on page 1149, “**posix\_spawn\_file\_actions\_addclose** or **posix\_spawn\_file\_actions\_addopen** Subroutine” on page 1153, “**posix\_spawn\_file\_actions\_adddup2** Subroutine” on page 1154, “**posix\_spawn\_file\_actions\_destroy** or **posix\_spawn\_file\_actions\_init** Subroutine” on page 1155, “**posix\_spawnattr\_destroy** or **posix\_spawnattr\_init** Subroutine” on page 1156, “**posix\_spawnattr\_getflags** or **posix\_spawnattr\_setflags** Subroutine” on page 1157, “**posix\_spawnattr\_getpgroup** or **posix\_spawnattr\_setpgroup** Subroutine” on page 1158, “**posix\_spawnattr\_getschedparam** or **posix\_spawnattr\_setschedparam** Subroutine” on page 1159, “**posix\_spawnattr\_getschedpolicy** or **posix\_spawnattr\_setschedpolicy** Subroutine” on page 1160, “**posix\_spawnattr\_getsigmask** or **posix\_spawnattr\_setsigmask** Subroutine.”

---

## **posix\_spawnattr\_getsigmask** or **posix\_spawnattr\_setsigmask** Subroutine

### Purpose

Gets and sets the **spawn-sigmask** attribute of a spawn attributes object.

### Syntax

```
#include <signal.h>
#include <spawn.h>

int posix_spawnattr_getsigmask(const posix_spawnattr_t *restrict attr,
                               sigset_t *restrict sigmask);
int posix_spawnattr_setsigmask(posix_spawnattr_t *restrict attr,
                               const sigset_t *restrict sigmask);
```

### Description

The **posix\_spawnattr\_getsigmask** subroutine gets the value of the **spawn-sigmask** attribute from the attributes object referenced by *attr*.

The **posix\_spawnattr\_setsigmask** subroutine sets the **spawn-sigmask** attribute in an initialized attributes object referenced by *attr*.

The **spawn-sigmask** attribute represents the signal mask in effect in the new process image of a spawn operation (if **POSIX\_SPAWN\_SETSIGMASK** is set in the **spawn-flags** attribute). The default value of this attribute is unspecified.

## Return Values

Upon successful completion, the **posix\_spawnattr\_getsigmask** subroutine returns 0 and stores the value of the **spawn-sigmask** attribute of *attr* into the object referenced by the *sigmask* parameter; otherwise, an error number is returned to indicate the error.

Upon successful completion, the **posix\_spawnattr\_setsigmask** subroutine returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The **posix\_spawnattr\_getsigmask** and **posix\_spawnattr\_setsigmask** subroutines might fail if the following is true:

**EINVAL**            The value specified by *attr* is invalid.

The **posix\_spawnattr\_setsigmask** subroutine might fail if the following is true:

**EINVAL**            The value of the attribute being set is not valid.

## Related Information

The “[posix\\_spawn or posix\\_spawn Subroutine](#)” on page 1149, “[posix\\_spawn\\_file\\_actions\\_addclose or posix\\_spawn\\_file\\_actions\\_addopen Subroutine](#)” on page 1153, “[posix\\_spawn\\_file\\_actions\\_adddup2 Subroutine](#)” on page 1154, “[posix\\_spawn\\_file\\_actions\\_destroy or posix\\_spawn\\_file\\_actions\\_init Subroutine](#)” on page 1155, “[posix\\_spawnattr\\_destroy or posix\\_spawnattr\\_init Subroutine](#)” on page 1156, “[posix\\_spawnattr\\_getsigdefault or posix\\_spawnattr\\_setsigdefault Subroutine](#)” on page 1161, “[posix\\_spawnattr\\_getflags or posix\\_spawnattr\\_setflags Subroutine](#)” on page 1157, “[posix\\_spawnattr\\_getpgroup or posix\\_spawnattr\\_setpgroup Subroutine](#)” on page 1158, “[posix\\_spawnattr\\_getschedparam or posix\\_spawnattr\\_setschedparam Subroutine](#)” on page 1159, “[posix\\_spawnattr\\_getschedpolicy or posix\\_spawnattr\\_setschedpolicy Subroutine](#)” on page 1160.

---

## posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event Subroutine

### Purpose

Retrieves a trace event (TRACING).

### Syntax

```
#include <sys/types.h>
#include <trace.h>

int posix_trace_getnext_event(trace_id_t trid,
    struct posix_trace_event_info *restrict event,
    void *restrict data, size_t num_bytes,
    size_t *restrict data_len, int *restrict unavailable);

int posix_trace_timedgetnext_event(trace_id_t trid,
    struct posix_trace_event_info *restrict event,
```

```
void *restrict data, size_t num_bytes,
size_t *restrict data_len, int *restrict unavailable,
const struct timespec *restrict abs_timeout);
```

```
int posix_trace_trygetnext_event(trace_id_t trid,
struct posix_trace_event_info *restrict event,
void *restrict data, size_t num_bytes,
size_t *restrict data_len, int *restrict unavailable);
```

## Description

The **posix\_trace\_getnext\_event()** function reports a recorded trace event either from an active trace stream without log or a prerecorded trace stream identified by the *trid* argument. The **posix\_trace\_trygetnext\_event()** function reports a recorded trace event from an active trace stream without log identified by the *trid* argument.

The trace event information associated with the recorded trace event is copied by the function into the structure pointed to by the *event* argument, and the data associated with the trace event is copied into the buffer pointed to by the *data* argument.

The **posix\_trace\_getnext\_event()** function blocks if the *trid* argument identifies an active trace stream and there is currently no trace event ready to be retrieved. When returning, if a recorded trace event was reported, the variable pointed to by the *unavailable* argument is set to 0. Otherwise, the variable pointed to by the *unavailable* argument is set to a value different from 0.

The **posix\_trace\_timedgetnext\_event()** function attempts to get another trace event from an active trace stream without log, as in the **posix\_trace\_getnext\_event()** function. However, if no trace event is available from the trace stream, the implied wait terminates when the timeout specified by the argument *abs\_timeout* expires, and the function returns the error [ETIMEDOUT].

The timeout expires when the absolute time specified by *abs\_timeout* passes—as measured by the clock upon which timeouts are based (that is, when the value of that clock equals or exceeds *abs\_timeout*)—or when the absolute time specified by *abs\_timeout* has already passed at the time of the call.

If the **Timers** option is supported, the timeout is based on the CLOCK\_REALTIME clock; if the **Timers** option is not supported, the timeout is based on the system clock as returned by the **time()** function. The resolution of the timeout matches the resolution of the clock on which it is based. The **timespec** data type is defined in the **<time.h>** header.

The function never fails with a timeout if a trace event is immediately available from the trace stream. The validity of the *abs\_timeout* argument does not need to be checked if a trace event is immediately available from the trace stream.

The behavior of this function for a prerecorded trace stream is unspecified.

The **posix\_trace\_trygetnext\_event()** function does not block. This function returns an error if the *trid* argument identifies a prerecorded trace stream. If a recorded trace event was reported, the variable pointed to by the *unavailable* argument is set to 0. Otherwise, if no trace event was reported, the variable pointed to by the *unavailable* argument is set to a value different from 0.

The *num\_bytes* argument equals the size of the buffer pointed to by the *data* argument. The *data\_len* argument reports to the application the length, in bytes, of the data record just transferred. If *num\_bytes* is greater than or equal to the size of the data associated with the trace event pointed to by the *event* argument, all the recorded data is transferred. In this case, the truncation-status member of the trace event structure is either POSIX\_TRACE\_NOT\_TRUNCATED (if the trace event data was recorded without truncation while tracing) or POSIX\_TRACE\_TRUNCATED\_RECORD (if the trace event data was truncated when it was recorded). If the *num\_bytes* argument is less than the length of recorded trace event data, the data transferred is truncated to a length of *num\_bytes*, the value stored in the variable pointed to by

*data\_len* equals *num\_bytes*, and the truncation-status member of the *event* structure argument is set to `POSIX_TRACE_TRUNCATED_READ` (see the `posix_trace_event_info` structure defined in `<trace.h>`).

The report of a trace event is sequential starting from the oldest recorded trace event. Trace events are reported in the order in which they were generated, up to an implementation-defined time resolution that causes the ordering of trace events occurring very close to each other to be unknown. After it is reported, a trace event cannot be reported again from an active trace stream. After a trace event is reported from an active trace stream without log, the trace stream makes the resources associated with that trace event available to record future generated trace events.

## Return Values

Upon successful completion, these functions return a value of 0. Otherwise, they return the corresponding error number.

If successful, these functions store:

- The recorded trace event in the object pointed to by *event*
- The trace event information associated with the recorded trace event in the object pointed to by *data*
- The length of this trace event information in the object pointed to by *data\_len*
- The value of 0 in the object pointed to by *unavailable*

## Error Codes

These functions fail if:

**EINVAL** The trace stream identifier argument *trid* is invalid.

The `posix_trace_getnext_event()` and `posix_trace_timedgetnext_event()` functions fail if:

**EINTR** The operation was interrupted by a signal, and so the call had no effect.

The `posix_trace_trygetnext_event()` function fails if:

**EINVAL** The trace stream identifier argument *trid* does not correspond to an active trace stream.

The `posix_trace_timedgetnext_event()` function fails if:

**EINVAL** There is no trace event immediately available from the trace stream, and the *timeout* argument is invalid.

**ETIMEDOUT** No trace event was available from the trace stream before the specified *timeout* expired.

## Related Information

“mq\_receive, mq\_timedreceive Subroutine” on page 870, “mq\_send, mq\_timedsend Subroutine” on page 872, “pthread\_mutex\_timedlock Subroutine” on page 1272, “pthread\_rwlock\_timedrdlock Subroutine” on page 1287, “pthread\_rwlock\_timedwrlock Subroutine” on page 1289.

The `sem_timedwait` subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The `pthread.h` and `time.h` files in *AIX 5L Version 5.3 Files Reference*.

---

## powf, powl, or pow Subroutine

### Purpose

Computes power.

### Syntax

```
#include <math.h>
```

```
float powf (x, y)
```

```
float x;
```

```
float y;
```

```
long double powl (x, y)
```

```
long double x, y;
```

```
double pow (x, y)
```

```
double x, y;
```

### Description

The **powf**, **powl**, and **pow** subroutines compute the value of  $x$  raised to the power  $y$ ,  $x^y$ . If  $x$  is negative, the application ensures that  $y$  is an integer value.

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$	Specifies the value of the base.
$y$	Specifies the value of the exponent.

### Return Values

Upon successful completion, the **pow**, **powf** and **powl** subroutines return the value of  $x$  raised to the power  $y$ .

For finite values of  $x < 0$ , and finite non-integer values of  $y$ , a domain error occurs and a NaN is returned.

If the correct value would cause overflow, a range error occurs and the **pow**, **powf**, and **powl** subroutines return **HUGE\_VAL**, **HUGE\_VALF**, and **HUGE\_VALL**, respectively.

If the correct value would cause underflow, and is not representable, a range error may occur, and 0.0 is returned.

If  $x$  or  $y$  is a NaN, a NaN is returned (unless specified elsewhere in this description).

For any value of  $y$  (including NaN), if  $x$  is  $+1$ , 1.0 is returned.

For any value of  $x$  (including NaN), if  $y$  is  $\pm 0$ , 1.0 is returned.

For any odd integer value of  $y > 0$ , if  $x$  is  $\pm 0$ ,  $\pm 0$  is returned.

For  $y > 0$  and not an odd integer, if  $x$  is  $\pm 0$ ,  $+0$  is returned.

If  $x$  is  $-1$ , and  $y$  is  $\pm\text{Inf}$ ,  $1.0$  is returned.

For  $|x| < 1$ , if  $y$  is  $-\text{Inf}$ ,  $+\text{Inf}$  is returned.

For  $|x| > 1$ , if  $y$  is  $-\text{Inf}$ ,  $+0$  is returned.

For  $|x| < 1$ , if  $y$  is  $+\text{Inf}$ ,  $+0$  is returned.

For  $|x| > 1$ , if  $y$  is  $+\text{Inf}$ ,  $+\text{Inf}$  is returned.

For  $y$  an odd integer  $< 0$ , if  $x$  is  $-\text{Inf}$ ,  $-0$  is returned.

For  $y < 0$  and not an odd integer, if  $x$  is  $-\text{Inf}$ ,  $+0$  is returned.

For  $y$  an odd integer  $> 0$ , if  $x$  is  $-\text{Inf}$ ,  $-\text{Inf}$  is returned.

For  $y > 0$  and not an odd integer, if  $x$  is  $-\text{Inf}$ ,  $+\text{Inf}$  is returned.

For  $y < 0$ , if  $x$  is  $+\text{Inf}$ ,  $+0$  is returned.

For  $y > 0$ , if  $x$  is  $+\text{Inf}$ ,  $+\text{Inf}$  is returned.

For  $y$  an odd integer  $< 0$ , if  $x$  is  $\pm 0$ , a pole error occurs and  $\pm\text{HUGE\_VAL}$ ,  $\pm\text{HUGE\_VALF}$ , and  $\pm\text{HUGE\_VALL}$  is returned for **pow**, **powf**, and **powl**, respectively.

For  $y < 0$  and not an odd integer, if  $x$  is  $\pm 0$ , a pole error occurs and **HUGE\_VAL**, **HUGE\_VALF** and **HUGE\_VALL** is returned for **pow**, **powf**, and **powl**, respectively.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

## Error Codes

When using the **libm.a** library:

- |             |   |
|-------------|---|
| <b>pow</b>  | If the correct value overflows, the <b>pow</b> subroutine returns a <b>HUGE_VAL</b> value and sets <b>errno</b> to <b>ERANGE</b> . If the $x$ parameter is negative and the $y$ parameter is not an integer, the <b>pow</b> subroutine returns a <b>NaNQ</b> value and sets <b>errno</b> to <b>EDOM</b> . If $x=0$ and the $y$ parameter is negative, the <b>pow</b> subroutine returns a <b>HUGE_VAL</b> value but does not modify <b>errno</b> .    |
| <b>powl</b> | If the correct value overflows, the <b>powl</b> subroutine returns a <b>HUGE_VAL</b> value and sets <b>errno</b> to <b>ERANGE</b> . If the $x$ parameter is negative and the $y$ parameter is not an integer, the <b>powl</b> subroutine returns a <b>NaNQ</b> value and sets <b>errno</b> to <b>EDOM</b> . If $x=0$ and the $y$ parameter is negative, the <b>powl</b> subroutine returns a <b>HUGE_VAL</b> value but does not modify <b>errno</b> . |

When using **libmsaa.a(-lmsaa)**:

**pow** If  $x=0$  and the  $y$  parameter is not positive, or if the  $x$  parameter is negative and the  $y$  parameter is not an integer, the **pow** subroutine returns 0 and sets **errno** to **EDOM**. In these cases a message indicating DOMAIN error is output to standard error. When the correct value for the **pow** subroutine would overflow or underflow, the **pow** subroutine returns:

**+HUGE\_VAL**

OR

**-HUGE\_VAL**

OR

0

When using either the **libm.a** library or the **libsaa.a** library:

**powl** If the correct value overflows, **powl** returns **HUGE\_VAL** and **errno** to **ERANGE**. If  $x$  is negative and  $y$  is not an integer, **powl** returns **NaNQ** and sets **errno** to **EDOM**. If  $x = 0$  and  $y$  is negative, **powl** returns a **HUGE\_VAL** value but does not modify **errno**.

## Related Information

“exp, expf, or expl Subroutine” on page 245, “feclearexcept Subroutine” on page 265, “fetetestexcept Subroutine” on page 273, and “class, \_class, finite, isnan, or unordered Subroutines” on page 167.

**math.h** in *AIX 5L Version 5.3 Files Reference*.

---

## printf, fprintf, sprintf, snprintf, wprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine

### Purpose

Prints formatted output.

### Library

Standard C Library (**libc.a**) or the Standard C Library with 128-Bit long doubles (**libc128.a**)

### Syntax

```
#include <stdio.h>
```

```
int printf (Format, [Value, ...])
const char *Format;
```

```
int fprintf (Stream, Format, [Value, ...])
FILE *Stream;
const char *Format;
```

```
int sprintf (String, Format, [Value, ...])
char *String;
const char *Format;
```

```
int snprintf (String, Number, Format, [Value, . . .])
char *String;
int Number;
const char *Format;
#include <stdarg.h>
```

```
int vprintf (Format, Value)
```



```

const char *Format;
va_list Value;

int vfprintf (Stream, Format, Value)
FILE *Stream;
const char *Format;
va_list Value;

int vsprintf (String, Format, Value)
char *String;
const char *Format;
va_list Value;
#include <wchar.h>

int wvsprintf (String, Format, Value)
wchar_t *String;
const char *Format;
va_list Value;

int wsprintf (String, Format, [Value, ...])
wchar_t *String;
const char *Format;

```

## Description

The **printf** subroutine converts, formats, and writes the *Value* parameter values, under control of the *Format* parameter, to the standard output stream. The **printf** subroutine provides conversion types to handle code points and **wchar\_t** wide character codes.

The **fprintf** subroutine converts, formats, and writes the *Value* parameter values, under control of the *Format* parameter, to the output stream specified by the *Stream* parameter. This subroutine provides conversion types to handle code points and **wchar\_t** wide character codes.

The **sprintf** subroutine converts, formats, and stores the *Value* parameter values, under control of the *Format* parameter, into consecutive bytes, starting at the address specified by the *String* parameter. The **sprintf** subroutine places a null character (\0) at the end. You must ensure that enough storage space is available to contain the formatted string. This subroutine provides conversion types to handle code points and **wchar\_t** wide character codes.

The **snprintf** subroutine converts, formats, and stores the *Value* parameter values, under control of the *Format* parameter, into consecutive bytes, starting at the address specified by the *String* parameter. The **snprintf** subroutine places a null character (\0) at the end. You must ensure that enough storage space is available to contain the formatted string. This subroutine provides conversion types to handle code points and **wchar\_t** wide character codes. The **snprintf** subroutine is identical to the **sprintf** subroutine with the addition of the *Number* parameter, which states the size of the buffer referred to by the *String* parameter.

The **wsprintf** subroutine converts, formats, and stores the *Value* parameter values, under control of the *Format* parameter, into consecutive **wchar\_t** characters starting at the address specified by the *String* parameter. The **wsprintf** subroutine places a null character (\0) at the end. The calling process should ensure that enough storage space is available to contain the formatted string. The field width unit is specified as the number of **wchar\_t** characters. The **wsprintf** subroutine is the same as the **printf** subroutine, except that the *String* parameter for the **wsprintf** subroutine uses a string of **wchar\_t** wide-character codes.

All of the above subroutines work by calling the **\_doprnt** subroutine, using variable-length argument facilities of the **varargs** macros.

The **vprintf**, **fprintf**, **vsprintf**, and **wvsprintf** subroutines format and write **varargs** macros parameter lists. These subroutines are the same as the **printf**, **fprintf**, **sprintf**, **snprintf**, and **wsprintf** subroutines,

respectively, except that they are not called with a variable number of parameters. Instead, they are called with a parameter-list pointer as defined by the **varargs** macros.

## Parameters

### *Number*

Specifies the number of bytes in a string to be copied or transformed.

*Value* Specifies 0 or more arguments that map directly to the objects in the *Format* parameter.

### *Stream*

Specifies the output stream.

*String* Specifies the starting address.

### *Format*

A character string that contains two types of objects:

- Plain characters, which are copied to the output stream.
- Conversion specifications, each of which causes 0 or more items to be retrieved from the *Value* parameter list. In the case of the **vprintf**, **vfprintf**, **vsprintf**, and **vwprintf** subroutines, each conversion specification causes 0 or more items to be retrieved from the **varargs** macros parameter lists.

If the *Value* parameter list does not contain enough items for the *Format* parameter, the results are unpredictable. If more parameters remain after the entire *Format* parameter has been processed, the subroutine ignores them.

Each conversion specification in the *Format* parameter has the following elements:

- A % (percent sign).
- 0 or more options, which modify the meaning of the conversion specification. The option characters and their meanings are:
  - ' Formats the integer portions resulting from **i**, **d**, **u**, **f**, **g** and **G** decimal conversions with **thousands\_sep** grouping characters. For other conversions the behavior is undefined. This option uses the nonmonetary grouping character.
  - Left-justifies the result of the conversion within the field.
  - + Begins the result of a signed conversion with a + (plus sign) or - (minus sign).

#### **space character**

Prefixes a space character to the result if the first character of a signed conversion is not a sign. If both the space-character and + option characters appear, the space-character option is ignored.

- # Converts the value to an alternate form. For **c**, **d**, **s**, and **u** conversions, the option has no effect. For **o** conversion, it increases the precision to force the first digit of the result to be a 0. For **x** and **X** conversions, a nonzero result has a 0x or 0X prefix. For **e**, **E**, **f**, **g**, and **G** conversions, the result always contains a decimal point, even if no digits follow it. For **g** and **G** conversions, trailing 0's are not removed from the result.
- 0 Pads to the field width with leading 0's (following any indication of sign or base) for **d**, **i**, **o**, **u**, **x**, **X**, **e**, **E**, **f**, **g**, and **G** conversions; the field is not space-padded. If the **0** and - options both appear, the **0** option is ignored. For **d**, **i**, **o**, **u**, **x**, and **X** conversions, if a precision is specified, the **0** option is also ignored. If the **0** and ' options both appear, grouping characters are inserted before the field is padded. For other conversions, the results are unreliable.
- B** Specifies a no-op character.
- N** Specifies a no-op character.
- J** Specifies a no-op character.

- An optional decimal digit string that specifies the minimum field width. If the converted value has fewer characters than the field width, the field is padded on the left to the length specified by the field width. If the - (left-justify) option is specified, the field is padded on the right.
- An optional precision. The precision is a . (dot) followed by a decimal digit string. If no precision is specified, the default value is 0. The precision specifies the following limits:
  - Minimum number of digits to appear for the **d**, **i**, **o**, **u**, **x**, or **X** conversions.
  - Number of digits to appear after the decimal point for the **e**, **E**, and **f** conversions.
  - Maximum number of significant digits for **g** and **G** conversions.
  - Maximum number of bytes to be printed from a string in **s** and **S** conversions.
  - Maximum number of bytes, converted from the **wchar\_t** array, to be printed from the **S** conversions. Only complete characters are printed.
- An optional **l** (lowercase *L*), **ll** (lowercase *LL*), **h**, or **L** specifier indicates one of the following:
  - An optional **h** specifying that a subsequent **d**, **i**, **u**, **o**, **x**, or **X** conversion specifier applies to a **short int** or **unsigned short int** *Value* parameter (the parameter will have been promoted according to the integral promotions, and its value will be converted to a **short int** or **unsigned short int** before printing).
  - An optional **h** specifying that a subsequent **n** conversion specifier applies to a pointer to a **short int** parameter.
  - An optional **l** (lowercase *L*) specifying that a subsequent **d**, **i**, **u**, **o**, **x**, or **X** conversion specifier applies to a **long int** or **unsigned long int** parameter .
  - An optional **l** (lowercase *L*) specifying that a subsequent **n** conversion specifier applies to a pointer to a **long int** parameter.
  - An optional **ll** (lowercase *LL*) specifying that a subsequent **d**, **i**, **u**, **o**, **x**, or **X** conversion specifier applies to a **long long int** or **unsigned long long int** parameter.
  - An optional **ll** (lowercase *LL*) specifying that a subsequent **n** conversion specifier applies to a pointer to a **long long int** parameter.
  - An optional **L** specifying that a following **e**, **E**, **f**, **g**, or **G** conversion specifier applies to a **long double** parameter. If linked with **libc.a**, **long double** is the same as double (64bits). If linked with **libc128.a** and **libc.a**, **long double** is 128 bits.
- An optional **H**, **D**, or **DD** specifier indicates one of the following conversions:
  - An optional **H** specifying that a following **e**, **E**, **f**, **F**, **g**, or **G** conversion specifier applies to a **\_Decimal32** parameter.
  - An optional **D** specifying that a following **e**, **E**, **f**, **F**, **g**, or **G** conversion specifier applies to a **\_Decimal64** parameter.
  - An optional **DD** specifying that a following **e**, **E**, **f**, **F**, **g**, or **G** conversion specifier applies to a **\_Decimal128** parameter.
- An optional **vl**, **lv**, **vh**, **hv** or **v** specifier indicates one of the following vector data type conversions:
  - An optional **v** specifying that a following **e**, **E**, **f**, **g**, **G**, **a**, or **A** conversion specifier applies to a **vector float** parameter. It consumes one argument and interprets the data as a series of four 4-byte floating point components.
  - An optional **v** specifying that a following **c**, **d**, **i**, **u**, **o**, **x**, or **X** conversion specifier applies to a **vector signed char**, **vector unsigned char**, or **vector bool char** parameter. It consumes one argument and interprets the data as a series of sixteen 1-byte components.
  - An optional **vl** or **lv** specifying that a following **d**, **i**, **u**, **o**, **x**, or **X** conversion specifier applies to a **vector signed int**, **vector unsigned int**, or **vector bool** parameter. It consumes one argument and interprets the data as a series of four 4-byte integer components.
  - An optional **vh** or **hv** specifying that a following **d**, **i**, **u**, **o**, **x**, or **X** conversion specifier applies to a **vector signed short** or **vector unsigned short** parameter. It consumes one argument and interprets the data as a series of eight 2-byte integer components.

- For any of the preceding specifiers, an optional separator character can be specified immediately preceding the vector size specifier. If no separator is specified, the default separator is a space unless the conversion is **c**, in which case the default separator is null. The set of supported optional separators are , (comma), ; (semicolon), : (colon), and \_ (underscore).
- The following characters indicate the type of conversion to be applied:
  - %** Performs no conversion. Prints (%).
  - d or i** Accepts a *Value* parameter specifying an integer and converts it to signed decimal notation. The precision specifies the minimum number of digits to appear. If the value being converted can be represented in fewer digits, it is expanded with leading 0's. The default precision is 1. The result of converting a value of 0 with a precision of 0 is a null string. Specifying a field width with a 0 as a leading character causes the field-width value to be padded with leading 0's.
  - u** Accepts a *Value* parameter specifying an unsigned integer and converts it to unsigned decimal notation. The precision specifies the minimum number of digits to appear. If the value being converted can be represented in fewer digits, it is expanded with leading 0's. The default precision is 1. The result of converting a value of 0 with a precision of 0 is a null string. Specifying a field width with a 0 as a leading character causes the field-width value to be padded with leading 0's.
  - o** Accepts a *Value* parameter specifying an unsigned integer and converts it to unsigned octal notation. The precision specifies the minimum number of digits to appear. If the value being converted can be represented in fewer digits, it is expanded with leading 0's. The default precision is 1. The result of converting a value of 0 with a precision of 0 is a null string. Specifying a field-width with a 0 as a leading character causes the field width value to be padded with leading 0's. An octal value for field width is not implied.
  - x or X** Accepts a *Value* parameter specifying an unsigned integer and converts it to unsigned hexadecimal notation. The letters **abcdef** are used for the **x** conversion and the letters **ABCDEF** are used for the **X** conversion. The precision specifies the minimum number of digits to appear. If the value being converted can be represented in fewer digits, it is expanded with leading 0's. The default precision is 1. The result of converting a value of 0 with a precision of 0 is a null string. Specifying a field width with a 0 as a leading character causes the field-width value to be padded with leading 0's.
  - f** Accepts a *Value* parameter specifying a double and converts it to decimal notation in the format [-]ddd.ddd. The number of digits after the decimal point is equal to the precision specification. If no precision is specified, six digits are output. If the precision is 0, no decimal point appears.
  - e or E** Accepts a *Value* parameter specifying a double and converts it to the exponential form [-]d.ddde+/-dd. One digit exists before the decimal point, and the number of digits after the decimal point is equal to the precision specification. The precision specification can be in the range of 0-17 digits. If no precision is specified, six digits are output. If the precision is 0, no decimal point appears. The **E** conversion character produces a number with **E** instead of **e** before the exponent. The exponent always contains at least two digits.
  - g or G** Accepts a *Value* parameter specifying a double and converts it in the style of the **e**, **E**, or **f** conversion characters, with the precision specifying the number of significant digits. Trailing 0's are removed from the result. A decimal point appears only if it is followed by a digit. The style used depends on the value converted. Style **e** (**E**, if **G** is the flag used) results only if the exponent resulting from the conversion is less than -4, or if it is greater or equal to the precision. If an explicit precision is 0, it is taken as 1.

- c** Accepts and prints a *Value* parameter specifying an integer converted to an **unsigned char** data type.
- C** Accepts and prints a *Value* parameter specifying a **wchar\_t** wide character code. The **wchar\_t** wide character code specified by the *Value* parameter is converted to an array of bytes representing a character and that character is written; the *Value* parameter is written without conversion when using the **wsprintf** subroutine.
- s** Accepts a *Value* parameter as a string (character pointer), and characters from the string are printed until a null character (\0) is encountered or the number of bytes indicated by the precision is reached. If no precision is specified, all bytes up to the first null character are printed. If the string pointer specified by the *Value* parameter has a null value, the results are unreliable.
- S** Accepts a corresponding *Value* parameter as a pointer to a **wchar\_t** string. Characters from the string are printed (without conversion) until a null character (\0) is encountered or the number of wide characters indicated by the precision is reached. If no precision is specified, all characters up to the first null character are printed. If the string pointer specified by the *Value* parameter has a value of null, the results are unreliable.
- p** Accepts a pointer to void. The value of the pointer is converted to a sequence of printable characters, the same as an unsigned hexadecimal (x).
- n** Accepts a pointer to an integer into which is written the number of characters (wide-character codes in the case of the **wsprintf** subroutine) written to the output stream by this call. No argument is converted.

A field width or precision can be indicated by an \* (asterisk) instead of a digit string. In this case, an integer *Value* parameter supplies the field width or precision. The *Value* parameter converted for output is not retrieved until the conversion letter is reached, so the parameters specifying field width or precision must appear before the value (if any) to be converted.

If the result of a conversion is wider than the field width, the field is expanded to contain the converted result and no truncation occurs. However, a small field width or precision can cause truncation on the right.

The **printf**, **fprintf**, **sprintf**, **snprintf**, **wsprintf**, **vprintf**, **vfprintf**, **vsprintf**, or **vwsprintf** subroutine allows the insertion of a language-dependent radix character in the output string. The radix character is defined by language-specific data in the **LC\_NUMERIC** category of the program's locale. In the C locale, or in a locale where the radix character is not defined, the radix character defaults to a . (dot).

After any of these subroutines runs successfully, and before the next successful completion of a call to the **fclose** ("fclose or fflush Subroutine" on page 253) or **fflush** subroutine on the same stream or to the **exit** ("exit, atexit, unatexit, \_exit, or \_Exit Subroutine" on page 243) or **abort** ("abort Subroutine" on page 2) subroutine, the *st\_ctime* and *st\_mtime* fields of the file are marked for update.

The **e**, **E**, **f**, **g**, and **G** conversion specifiers represent the special floating-point values as follows:

<b>Quiet NaN</b>	+NaNQ or -NaNQ
<b>Signaling NaN</b>	+NaNS or -NaNS
<b>+/-INF</b>	+INF or -INF
<b>+/-0</b>	+0 or -0

The representation of the + (plus sign) depends on whether the + or space-character formatting option is specified.

These subroutines can handle a format string that enables the system to process elements of the parameter list in variable order. In such a case, the normal conversion character % (percent sign) is replaced by *%digit\$*, where *digit* is a decimal number in the range from 1 to the **NL\_ARGMAX** value. Conversion is then applied to the specified argument, rather than to the next unused argument. This feature provides for the definition of format strings in an order appropriate to specific languages. When variable ordering is used the \* (asterisk) specification for field width in precision is replaced by *%digit\$*. If you use the variable-ordering feature, you must specify it for all conversions.

The following criteria apply:

- The format passed to the NLS extensions can contain either the format of the conversion or the explicit or implicit argument number. However, these forms cannot be mixed within a single format string, except for %% (double percent sign).
- The *n* value must have no leading zeros.
- If *%n\$* is used, *%1\$* to *%n - 1\$* inclusive must be used.
- The *n* in *%n\$* is in the range from 1 to the **NL\_ARGMAX** value, inclusive. See the **limits.h** file for more information about the **NL\_ARGMAX** value.
- Numbered arguments in the argument list can be referenced as many times as required.
- The \* (asterisk) specification for field width or precision is not permitted with the variable order *%n\$* format; instead, the *\*m\$* format is used.

## Return Values

Upon successful completion, the **printf**, **fprintf**, **vprintf**, and **vfprintf** subroutines return the number of bytes transmitted (not including the null character [0] in the case of the **sprintf**, and **vsprintf** subroutines). If an error was encountered, a negative value is output.

Upon successful completion, the **snprintf** subroutine returns the number of bytes written to the *String* parameter (excluding the terminating null byte). If output characters are discarded because the output exceeded the *Number* parameter in length, then the **snprintf** subroutine returns the number of bytes that would have been written to the *String* parameter if the *Number* parameter had been large enough (excluding the terminating null byte).

Upon successful completion, the **wsprintf** and **vwsprintf** subroutines return the number of wide characters transmitted (not including the wide character null character [0]). If an error was encountered, a negative value is output.

## Error Codes

The **printf**, **fprintf**, **sprintf**, **snprintf**, or **wsprintf** subroutine is unsuccessful if the file specified by the *Stream* parameter is unbuffered or the buffer needs to be flushed and one or more of the following are true:

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> or <b>O_NDELAY</b> flag is set for the file descriptor underlying the file specified by the <i>Stream</i> or <i>String</i> parameter and the process would be delayed in the write operation.
<b>EBADF</b>	The file descriptor underlying the file specified by the <i>Stream</i> or <i>String</i> parameter is not a valid file descriptor open for writing.
<b>EFBIG</b>	An attempt was made to write to a file that exceeds the file size limit of this process or the maximum file size. For more information, refer to the <b>ulimit</b> subroutine.
<b>EINTR</b>	The write operation terminated due to receipt of a signal, and either no data was transferred or a partial transfer was not reported.

**Note:** Depending upon which library routine the application binds to, this subroutine may return **EINTR**. Refer to the **signal** subroutine regarding **sa\_restart**.

- EIO** The process is a member of a background process group attempting to perform a write to its controlling terminal, the **TOSTOP** flag is set, the process is neither ignoring nor blocking the **SIGTTOU** signal, and the process group of the process has no parent process.
- ENOSPC** No free space remains on the device that contains the file.
- EPIPE** An attempt was made to write to a pipe or first-in-first-out (FIFO) that is not open for reading by any process. A **SIGPIPE** signal is sent to the process.

The **printf**, **fprintf**, **sprintf**, **snprintf**, or **wsprintf** subroutine may be unsuccessful if one or more of the following are true:

- EILSEQ** An invalid character sequence was detected.
- EINVAL** The *Format* parameter received insufficient arguments.
- ENOMEM** Insufficient storage space is available.
- ENXIO** A request was made of a nonexistent device, or the request was outside the capabilities of the device.

## Examples

The following example demonstrates how the **vfprintf** subroutine can be used to write an error routine:

```
#include <stdio.h>
#include <stdarg.h>
/* The error routine should be called with the
   syntax:          */
/* error(routine_name, Format
   [, value, . . . ]); */
/*VARARGS0*/
void error(char *fmt, . . .);
/* ** Note that the function name and
   Format arguments cannot be **
   separately declared because of the **
   definition of varargs. */ {
   va_list args;

   va_start(args, fmt);
   /*
   ** Display the name of the function
   that called the error routine */
   fprintf(stderr, "ERROR in %s: ",
           va_arg(args, char *)); /*
   ** Display the remainder of the message
   */
   fmt = va_arg(args, char *);
   vfprintf(fmt, args);
   va_end(args);
   abort(); }
```

## Related Information

The **abort** (“abort Subroutine” on page 2) subroutine, **conv** (“conv Subroutines” on page 183) subroutine, **ecvt**, **fcvt**, or **gcvt** (“ecvt, fcvt, or gcvt Subroutine” on page 225) subroutine, **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine, **fclose** or **fflush** (“fclose or fflush Subroutine” on page 253) subroutine, **putc**, **putchar**, **fputc**, or **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **putwc**, **putwchar**, or **fputwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **scanf**, **fscanf**, **sscanf**, or **wscanf** subroutine, **setlocale** subroutine.

Input and Output Handling and 128-Bit Long Double Floating-Point Data Type in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## profil Subroutine

### Purpose

Starts and stops program address sampling for execution profiling.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <mon.h>
```

```
void profil ( ShortBuffer, BufferSize, Offset, Scale)
```

OR

```
void profil ( ProfBuffer, -1, 0, 0)
```

```
unsigned short *ShortBuffer;  
struct prof *ProfBuffer;  
unsigned int BufferSize, Scale;  
unsigned long Offset;
```

### Description

The **profil** subroutine arranges to record a histogram of periodically sampled values of the calling process program counter. If *BufferSize* is not -1:

- The parameters to the **profil** subroutine are interpreted as shown in the first syntax definition.
- After this call, the program counter (pc) of the process is examined each clock tick if the process is the currently active process. The value of the *Offset* parameter is subtracted from the pc. The result is multiplied by the value of the *Scale* parameter, shifted right 16 bits, and rounded up to the next half-word aligned value. If the resulting number is less than the *BufferSize* value divided by **sizeof(short)**, the corresponding **short** inside the *ShortBuffer* parameter is incremented. If the result of this increment would overflow an unsigned short, it remains USHRT\_MAX.
- The least significant 16 bits of the *Scale* parameter are interpreted as an unsigned, fixed-point fraction with a binary point at the left. The most significant 16 bits of the *Scale* parameter are ignored. For example:

Octal	Hex	Meaning
0177777	0xFFFF	Maps approximately each pair of bytes in the instruction space to a unique <b>short</b> in the <i>ShortBuffer</i> parameter.
077777	0x7FFF	Maps approximately every four bytes to a <b>short</b> in the <i>ShortBuffer</i> parameter.
02	0x0002	Maps all instructions to the same location, producing a noninterrupting core clock.
01	0x0001	Turns profiling off.
00	0x0000	Turns profiling off.

**Note:** Mapping each byte of the instruction space to an individual **short** in the *ShortBuffer* parameter is not possible.

- Profiling, using the first syntax definition, is rendered ineffective by giving a value of 0 for the *BufferSize* parameter.

If the value of the *BufferSize* parameter is -1:



- The parameters to the **profil** subroutine are interpreted as shown in the second syntax definition. In this case, the *Offset* and *Scale* parameters are ignored, and the *ProfBuffer* parameter points to an array of **prof** structures. The **prof** structure is defined in the **mon.h** file, and it contains the following members:

```

caddr_t      p_low;
caddr_t      p_high;
HISTCOUNTER  *p_buff;
int          p_bufsize;
uint         p_scale;

```

If the *p\_scale* member has the value of -1, a value for it is computed based on *p\_low*, *p\_high*, and *p\_bufsize*; otherwise *p\_scale* is interpreted like the *scale* argument in the first synopsis. The *p\_high* members in successive structures must be in ascending sequence. The array of structures is ended with a structure containing a *p\_high* member set to 0; all other fields in this last structure are ignored.

The *p\_buff* buffer pointers in the array of **prof** structures must point into a single contiguous buffer space.

- Profiling, using the second syntax definition, is turned off by giving a *ProfBuffer* argument such that the *p\_high* element of the first structure is equal to 0.

In every case:

- Profiling remains on in both the child process and the parent process after a **fork** subroutine.
- Profiling is turned off when an **exec** subroutine is run.
- A call to the **profil** subroutine is ineffective if profiling has been previously turned on using one syntax definition, and an attempt is made to turn profiling off using the other syntax definition.
- A call to the **profil** subroutine is ineffective if the call is attempting to turn on profiling when profiling is already turned on, or if the call is attempting to turn off profiling when profiling is already turned off.

## Parameters

<i>ShortBuffer</i>	Points to an area of memory in the user address space. Its length (in bytes) is given by the <i>BufferSize</i> parameter.
<i>BufferSize</i>	Specifies the length (in bytes) of the buffer.
<i>Offset</i>	Specifies the delta of program counter start and buffer; for example, a 0 <i>Offset</i> implies that text begins at 0. If the user wants to use the entry point of a routine for the <i>Offset</i> parameter, the syntax of the parameter is as follows:  <b>*(long *)RoutineName</b>
<i>Scale</i>	Specifies the mapping factor between the program counter and <i>ShortBuffer</i> .
<i>ProfBuffer</i>	Points to an array of <b>prof</b> structures.

## Return Values

The **profil** subroutine always returns a value of 0. Otherwise, the **errno** global variable is set to indicate the error.

## Error Codes

The **profil** subroutine is unsuccessful if one or both of the following are true:

- |               |  |
|---------------|--|
| <b>EFAULT</b> | The address specified by the <i>ShortBuffer</i> or <i>ProfBuffer</i> parameters is not valid, or the address specified by a <i>p_buff</i> field is not valid. EFAULT can also occur if there are not sufficient resources to pin the profiling buffer in real storage. |
| <b>EINVAL</b> | The <i>p_high</i> fields in the <b>prof</b> structure specified by the <i>ProfBuffer</i> parameter are not in ascending order.   |

## Related Information

The **exec** (“exec: execl, execlp, execl, execlp, execlp, execlp, or execlp Subroutine” on page 236) subroutines, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **moncontrol** (“moncontrol Subroutine” on page 849) subroutine, **monitor** (“monitor Subroutine” on page 851) subroutine, **monstartup** (“monstartup Subroutine” on page 856) subroutine.

The **prof** command.

---

## proj\_execve Subroutine

### Purpose

Executes an application with the specified project assignment.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
int proj_execve(char * path char *const arg[], char *const env[], projid_t projid, int force);
```

### Description

The **proj\_execve** system call assigns the requested project ID to the calling process and runs the given program. This subroutine checks whether the caller is allowed to assign the requested project ID to the application, using the available project assignment rules for the caller’s user ID, group ID, and application name. If the requested project assignment is not allowed, an error code is returned. However, the user with root authority or advanced accounting administrator capabilities can force the project assignment by setting the *force* parameter to 1.

### Parameters

<i>path</i>	Path for the application or program to be run.
<i>arg</i>	List of arguments for the new process.
<i>env</i>	Environment for the new process.
<i>projid</i>	Project ID to be assigned to the new process.
<i>force</i>	Option to override the allowed project list for the application, user, or group.

### Return Values

0	Upon success, does not return to the calling process.
-1	The subroutine failed.

### Error Codes

<b>EPERM</b>	Permission denied. A user without privileges attempted the call.
--------------	--

## Related Information

The “addproj Subroutine” on page 31, “chprojattr Subroutine” on page 158, “getproj Subroutine” on page 418, rmproj Subroutine.

## projdballoc Subroutine

### Purpose

Allocates a project database handle.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
projdballoc(void **handle)
```

### Description

The **projdballoc** subroutine allocates a handle to operate on the project database. By default, this *handle* is initialized to operate on the system project database; however, it can be reset with the **projdbfinit** subroutine to reference another project database.

### Parameters

*handle*            Pointer to a void pointer

### Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

### Return Values

0	Success
-1	Failure

### Error Codes

<b>EINVAL</b>	The passed pointer is NULL
<b>ENOMEM</b>	No space left on memory

### Related Information

The “addprojdb Subroutine” on page 32, “chprojattrdb Subroutine” on page 159, “getfirstprojdb Subroutine” on page 368, “getnextprojdb Subroutine” on page 396, “getprojdb Subroutine” on page 419, “projdbfinit Subroutine,” “projdbfree Subroutine” on page 1181, rmprojdb Subroutine.

---

## projdbfinit Subroutine

### Purpose

Sets the handle to use a local project database as specified in the dbfile pointer and opens the file with the specified mode.

## Library

The **libaacct.a** library.

## Syntax

<sys/aacct.h>

```
projdbfinit(void *handle, char *file, int mode)
```

## Description

The **projdbfinit** subroutine sets the specified *handle* to use the specified project definition file. The file is opened in the specified mode. Subsequently, the project database, as represented by the *handle* parameter, will be referenced through file system primitives.

The project database must be initialized before calling this subroutine. The routines **projdballoc** and **projdbfinit** are provided for this purpose. The specified file is opened in the specified mode. File system calls are used to operate on these types of files. The struct **projdb** is filled as follows:

```
projdb.type = PROJ_LOCAL
```

```
projdb.fdes = value returned from open() call.
```

If the *file* parameter is NULL, then the system project database is opened.

## Parameters

<i>handle</i>	Pointer to handle
<i>file</i>	Indicate the project definition file name
<i>mode</i>	Indicates the mode in which the file is opened

## Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

## Return Values

0	Success
-1	Failure

## Error Codes

<b>EINVAL</b>	Passed handle or file is invalid
---------------	----------------------------------

## Related Information

The “addprojdb Subroutine” on page 32, “chprojattrdb Subroutine” on page 159, “getfirstprojdb Subroutine” on page 368, “getnextprojdb Subroutine” on page 396, “getproj Subroutine” on page 418, “getprojdb Subroutine” on page 419, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, “projdbfree Subroutine” on page 1181, rmprojdb Subroutine.

---

## projdbfree Subroutine

### Purpose

Frees an allocated project database handle.

### Library

The **libaacct.a** library.

### Syntax

<sys/aacct.h>

```
projdbfree(void *handle)
```

### Description

The **projdbfree** subroutine releases the memory allocated to a project database handle. The closure operation is based on the type of project database. If a project database is local, then it is closed using system primitives. The project database must be initialized before calling this subroutine. The routines **projdballoc** and **projdbfinit** are provided for this purpose.

### Parameters

*handle*            Pointer to a void pointer

### Security

Only for privileged users. Privilege can be extended to nonroot users by granting the CAP\_AACCT capability to a user.

### Return Values

0	Success
-1	Failure

### Error Codes

**EINVAL**            Passed pointer is NULL

### Related Information

The “addprojdb Subroutine” on page 32, “chprojattrdb Subroutine” on page 159, “getfirstprojdb Subroutine” on page 368, “getnextprojdb Subroutine” on page 396, “getproj Subroutine” on page 418, “getprojdb Subroutine” on page 419, “projdballoc Subroutine” on page 1179, “projdbfinit Subroutine” on page 1179, rmpojdb Subroutine.

---

## psdanger Subroutine

### Purpose

Defines the amount of free paging space available.

## Syntax

```
#include <signal.h>
#include <sys/vminfo.h>
```

```
blkcnt_t psdanger (Signal)
int Signal;
```

## Description

The **psdanger** subroutine returns the difference between the current number of free paging-space blocks and the paging-space thresholds of the system.

## Parameters

*Signal* Defines the signal.

## Return Values

If the value of the *Signal* parameter is 0, the return value is the total number of paging-space blocks defined in the system.

If the value of the *Signal* parameter is -1, the return value is the number of free paging-space blocks available in the system.

If the value of the *Signal* parameter is **SIGDANGER**, the return value is the difference between the current number of free paging-space blocks and the paging-space warning threshold. If the number of free paging-space blocks is less than the paging-space warning threshold, the return value is negative.

If the value of the *Signal* parameter is **SIGKILL**, the return value is the difference between the current number of free paging-space blocks and the paging-space kill threshold. If the number of free paging-space blocks is less than the paging-space kill threshold, the return value is negative.

## Related Information

The **swapoff** subroutine, **swapon** subroutine, **swapqry** subroutine.

The **chps** command, **lsps** command, **mkps** command, **rmps** command, **swapoff** command, **swapon** command.

Paging space in *Operating system and device management*.

Subroutines Overview and Understanding Paging Space Programming Requirements in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## psignal Subroutine or sys\_siglist Vector

### Purpose

Prints system signal messages.

### Library

Standard C Library (**libc.a**)

## Syntax

```
psignal ( Signal, String)
unsigned Signal;
char *String;
char *sys_siglist[ ];
```

## Description

The **psignal** subroutine produces a short message on the standard error file describing the indicated signal. First the *String* parameter is printed, then the name of the signal and a new-line character.

To simplify variant formatting of signal names, the **sys\_siglist** vector of message strings is provided. The signal number can be used as an index in this table to get the signal name without the new-line character. The **NSIG** defined in the **signal.h** file is the number of messages provided for in the table. It should be checked because new signals may be added to the system before they are added to the table.

## Parameters

*Signal* Specifies a signal. The signal number should be among those found in the **signal.h** file.  
*String* Specifies a string that is printed. Most usefully, the *String* parameter is the name of the program that incurred the signal.

## Related Information

The **pperror** (“pperror Subroutine” on page 1024) subroutine, **sigvec** subroutine.

---

## pthdb\_attr, pthdb\_cond, pthdb\_condattr, pthdb\_key, pthdb\_mutex, pthdb\_mutexattr, pthdb\_pthread, pthdb\_pthread\_key, pthdb\_rwlock, or pthdb\_rwlockattr Subroutine

### Purpose

Reports the pthread library objects.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>

int pthdb_pthread (pthdb_session_t session,
                  pthdb_pthread_t * pthreadp,
                  int cmd)
int pthdb_pthread_key(pthdb_session_t session,
                    pthread_key_t * keyp,
                    int cmd)
int pthdb_attr(pthdb_session_t session,
              pthdb_attr_t * attrp,
              int cmd)
int pthdb_cond (pthdb_session_t session,
                pthdb_cond_t * condp,
                int cmd)
int pthdb_condattr (pthdb_session_t session,
```

```

        pthread_condattr_t * condattrp,
        int cmd)
int pthread_key(pthread_session_t session,
               pthread_t pthread,
               pthread_key_t * keyp,
               int cmd)
int pthread_mutex(pthread_session_t session,
                 pthread_mutex_t * mutexp,
                 int cmd)
int pthread_mutexattr(pthread_session_t session,
                     pthread_mutexattr_t * mutexattrp,
                     int cmd)
int pthread_rwlock(pthread_session_t session,
                  pthread_rwlock_t * rwlockp,
                  int cmd)
int pthread_rwlockattr(pthread_session_t session,
                      pthread_rwlockattr_t * rwlockattrp,
                      int cmd)

```

## Description

The pthread library maintains internal lists of objects: pthreads, mutexes, mutex attributes, condition variables, condition variable attributes, read/write locks, read/write lock attributes, attributes, pthread specific keys, and active keys. The pthread debug library provides access to these lists one element at a time via the functions listed above.

Each one of those functions acquire the next element in the list of objects. For example, the **pthread\_attr** function gets the next attribute on the list of attributes.

A report of **PTHDB\_INVALID\_OBJECT** represents the empty list or the end of a list, where *OBJECT* is equal to **PTHREAD**, **ATTR**, **MUTEX**, **MUTEXATTR**, **COND**, **CONDATTR**, **RWLOCK**, **RWLOCKATTR**, **KEY**, or **TID** as appropriate.

Each list is reset to the top of the list when the **pthread\_session\_update** function is called, or when the list function reports a **PTHDB\_INVALID\_\*** value. For example, when **pthread\_attr** reports an attribute of **PTHDB\_INVALID\_ATTR** the list is reset to the beginning such that the next call reports the first attribute in the list, if any.

When **PTHDB\_LIST\_FIRST** is passed for the *cmd* parameter, the first item in the list is retrieved.

## Parameters

<i>session</i>	Session handle.
<i>attrp</i>	Attribute object.
<i>cmd</i>	Reset to the beginning of the list.
<i>condp</i>	Pointer to Condition variable object.
<i>condattrp</i>	Pointer to Condition variable attribute object.
<i>keyp</i>	Pointer to Key object.
<i>mutexattrp</i>	Pointer to Mutex attribute object.
<i>mutexp</i>	Pointer to Mutex object.
<i>pthread</i>	pthread object.
<i>pthreadp</i>	Pointer to pthread object.
<i>rwlockp</i>	Pointer to Read/Write lock object.
<i>rwlockattrp</i>	Pointer to Read/Write lock attribute object.



## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_CMD</b>	Invalid command.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_INTERNAL</b>	Error in library.
<b>PTHDB_MEMORY</b>	Not enough memory

## Related Information

The **pthdebug.h** file.

The **pthread.h** file.

---

**pthdb\_attr\_detachstate, pthdb\_attr\_addr,  
pthdb\_attr\_guardsize, pthdb\_attr\_inheritsched,  
pthdb\_attr\_schedparam, pthdb\_attr\_schedpolicy,  
pthdb\_attr\_schedpriority, pthdb\_attr\_scope,  
pthdb\_attr\_stackaddr, pthdb\_attr\_stacksize, or pthdb\_attr\_suspendstate**  
**Subroutine**

## Purpose

Query the various fields of a pthread attribute and return the results in the specified buffer.

## Library

pthread debug library (**libpthdebug.a**)

## Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_attr_detachstate (pthdb_session_t session,  
                           pthdb_attr_t attr,  
                           pthdb_detachstate_t * detachstatep);  
  
int pthdb_attr_addr (pthdb_session_t session,  
                    pthdb_attr_t attr,  
                    pthdb_addr_t * addrp);  
  
int pthdb_attr_guardsize (pthdb_session_t session,  
                          pthdb_attr_t attr,  
                          pthdb_size_t * guardsizep);  
  
int pthdb_attr_inheritsched (pthdb_session_t session,  
                              pthdb_attr_t attr,  
                              pthdb_inheritsched_t * inheritschedp);  
  
int pthdb_attr_schedparam (pthdb_session_t session,  
                            pthdb_attr_t attr,  
                            struct sched_param * schedparamp);  
  
int pthdb_attr_schedpolicy (pthdb_session_t session,  
                             pthdb_attr_t attr,  
                             pthdb_policy_t * schedpolicyp)
```

```

int pthread_attr_schedpriority (pthread_session_t session,
                               pthread_attr_t attr,
                               int * schedpriorityp)

int pthread_attr_scope (pthread_session_t session,
                       pthread_attr_t attr,
                       pthread_scope_t * scopep)

int pthread_attr_stackaddr (pthread_session_t session,
                            pthread_attr_t attr,
                            pthread_size_t * stackaddrp);

int pthread_attr_stacksize (pthread_session_t session,
                            pthread_attr_t attr,
                            pthread_size_t * stacksizep);

int pthread_attr_suspendstate (pthread_session_t session,
                               pthread_attr_t attr,
                               pthread_suspendstate_t * suspendstatep)

```

## Description

Each pthread is created using either the default pthread attribute or a user-specified pthread attribute. These functions query the various fields of a pthread attribute and, if successful, return the result in the buffer specified. In all cases, the values returned reflect the expected fields of a pthread created with the attribute specified.

**pthread\_attr\_detachstate** reports if the created pthread is detachable (**PDS\_DETACHED**) or joinable (**PDS\_JOINABLE**). **PDS\_NOTSUP** is reserved for unexpected results.

**pthread\_attr\_addr** reports the address of the pthread\_attr\_t.

**pthread\_attr\_guardsize** reports the guard size for the attribute.

**pthread\_attr\_inheritsched** reports whether the created pthread will run with scheduling policy and scheduling parameters from the created pthread (**PIS\_INHERIT**), or from the attribute (**PIS\_EXPLICIT**). **PIS\_NOTSUP** is reserved for unexpected results.

**pthread\_attr\_schedparam** reports the scheduling parameters associated with the pthread attribute. See **pthread\_attr\_inheritsched** for additional information.

**pthread\_attr\_schedpolicy** reports whether the scheduling policy associated with the pthread attribute is other (**SP\_OTHER**), first in first out (**SP\_FIFO**), or round robin (**SP\_RR**). **SP\_NOTSUP** is reserved for unexpected results.

**pthread\_attr\_schedpriority** reports the scheduling priority associated with the pthread attribute. See **pthread\_attr\_inheritsched** for additional information.

**pthread\_attr\_scope** reports whether the created pthread will have process scope (**PS\_PROCESS**) or system scope (**PS\_SYSTEM**). **PS\_NOTSUP** is reserved for unexpected results.

**pthread\_attr\_stackaddr** reports the address of the stack.

**pthread\_attr\_stacksize** reports the size of the stack.

**pthread\_attr\_suspendstate** reports whether the created pthread will be suspended (**PSS\_SUSPENDED**) or not (**PSS\_UNSPENDED**). **PSS\_NOTSUP** is reserved for unexpected results.

## Parameters

*addr*

Attributes address.

<i>attr</i>	Attributes handle.
<i>detachstatep</i>	Detach state buffer.
<i>guardsizep</i>	Attribute guard size.
<i>inheritschedp</i>	Inherit scheduling buffer.
<i>schedparamp</i>	Scheduling parameters buffer.
<i>schedpolicyp</i>	Scheduling policy buffer.
<i>schedpriorityp</i>	Scheduling priority buffer.
<i>scopep</i>	Contention scope buffer.
<i>session</i>	Session handle.
<i>stackaddrp</i>	Attributes stack address.
<i>stacksizep</i>	Attributes stack size.
<i>suspendstatep</i>	Suspend state buffer.

## Return Values

If successful these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_ATTR</b>	Invalid attribute handle.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_NOTSUP</b>	Not supported.
<b>PTHDB_INTERNAL</b>	Internal library error.

## Related Information

The **pthdebug.h** file.

The **pthread.h** file.

---

## pthdb\_condattr\_pshared, or pthdb\_condattr\_addr Subroutine

### Purpose

Gets the condition variable attribute pshared value.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_condattr_pshared (pthdb_session_t session,
                           pthdb_condattr_t condattr,
                           pthdb_pshared_t * psharedp)
```

```
int pthdb_condattr_addr (pthdb_session_t session,
                        pthdb_condattr_t condattr,
                        pthdb_addr_t * addrp)
```

## Description

The `pthdb_condattr_pshared` function is used to get the condition variable attribute process shared value. The pshared value can be `PSH_SHARED`, `PSH_PRIVATE`, or `PSH_NOTSUP`.

The `pthdb_condattr_addr` function reports the address of the `pthread_condattr_t`.

## Parameters

<i>addrp</i>	Pointer to the address of the <code>pthread_condattr_t</code> .
<i>condattr</i>	Condition variable attribute handle
<i>psharedp</i>	Pointer to the pshared value.
<i>session</i>	Session handle.

## Return Values

If successful this function returns `PTHDB_SUCCESS`. Otherwise, an error code is returned.

## Error Codes

<code>PTHDB_BAD_CONDATTR</code>	Invalid condition variable attribute handle.
<code>PTHDB_BAD_SESSION</code>	Invalid session handle.
<code>PTHDB_CALLBACK</code>	Debugger call back error.
<code>PTHDB_INTERNAL</code>	Error in library.
<code>PTHDB_POINTER</code>	Invalid pointer

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

## `pthdb_cond_addr`, `pthdb_cond_mutex` or `pthdb_cond_pshared` Subroutine

### Purpose

Gets the condition variable's mutex handle and pshared value.

### Library

pthread debug library (`libpthdebug.a`)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_cond_addr (pthdb_session_t session,  
                    pthdb_cond_t cond,  
                    pthdb_addr_t * addrp)
```

```
int pthdb_cond_mutex (pthdb_session_t session,  
                     pthdb_cond_t cond,  
                     pthdb_mutex_t * mutexp)
```

```
int pthdb_cond_pshared (pthdb_session_t session,
                       pthdb_cond_t cond,
                       pthdb_pshared_t * psharedp)
```

## Description

The **pthdb\_cond\_addr** function reports the address of the `pthdb_cond_t`.

The **pthdb\_cond\_mutex** function is used to get the mutex handle associated with the particular condition variable, if the mutex does not exist then `PTHDB_INVALID_MUTEX` is returned from the mutex.

The **pthdb\_cond\_pshared** function is used to get the condition variable process shared value. The `pshared` value can be `PSH_SHARED`, `PSH_PRIVATE`, or `PSH_NOTSUP`.

## Parameters

<i>addr</i>	Condition variable address
<i>cond</i>	Condition variable handle
<i>mutexp</i>	Pointer to mutex
<i>psharedp</i>	Pointer to pshared value
<i>session</i>	Session handle.

## Return Values

If successful, these functions return `PTHDB_SUCCESS`. Otherwise, an error code is returned.

## Error Codes

<code>PTHDB_BAD_COND</code>	Invalid cond handle.
<code>PTHDB_BAD_SESSION</code>	Invalid session handle.
<code>PTHDB_CALLBACK</code>	Debugger call back error.
<code>PTHDB_INVALID_MUTEX</code>	Invalid mutex.
<code>PTHDB_INTERNAL</code>	Error in library.
<code>PTHDB_POINTER</code>	Invalid pointer

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

## **pthdb\_mutexattr\_addr, pthdb\_mutexattr\_prioceiling, pthdb\_mutexattr\_protocol, pthdb\_mutexattr\_pshared or pthdb\_mutexattr\_type Subroutine**

### Purpose

Gets the mutex attribute `pshared`, priority ceiling, protocol, and type values.

### Library

pthread debug library (`libpthdebug.a`)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthread_mutexattr_addr (pthread_session_t session,
                           pthread_mutexattr_t mutexattr,
                           pthread_addr_t * addrp)
```

```
int pthread_mutexattr_protocol (pthread_session_t session,
                                pthread_mutexattr_t mutexattr,
                                pthread_protocol_t * protocolp)
```

```
int pthread_mutexattr_pshared (pthread_session_t session,
                               pthread_mutexattr_t mutexattr,
                               pthread_pshared_t * psharedp)
```

```
int pthread_mutexattr_type (pthread_session_t session,
                            pthread_mutexattr_t mutexattr,
                            pthread_mutex_type_t * typep)
```

## Description

The **pthread\_mutexattr\_addr** function reports the address of the `pthread_mutexattr_t`.

The **pthread\_mutexattr\_prioceiling** function is used to get the mutex attribute priority ceiling value.

The **pthread\_mutexattr\_protocol** function is used to get the mutex attribute protocol value. The protocol value can be **MP\_INHERIT**, **MP\_PROTECT**, **MP\_NONE**, or **MP\_NOTSUP**.

The **pthread\_mutexattr\_pshared** function is used to get the mutex attribute process shared value. The pshared value can be **PSH\_SHARED**, **PSH\_PRIVATE**, or **PSH\_NOTSUP**.

The **pthread\_mutexattr\_type** is used to get the value of the mutex attribute type. The values for the mutex type can be **MK\_NONRECURSIVE\_NP**, **MK\_RECURSIVE\_NP**, **MK\_FAST\_NP**, **MK\_ERRORCHECK**, **MK\_RECURSIVE**, **MK\_NORMAL**, or **MK\_NOTSUP**.

## Parameters

<i>addr</i>	Mutex attribute address.
<i>mutexattr</i>	Condition variable attribute handle
<i>prioceiling</i>	Pointer to priority ceiling value.
<i>protocolp</i>	Pointer to protocol value.
<i>psharedp</i>	Pointer to pshared value.
<i>session</i>	Session handle.
<i>typep</i>	Pointer to type value.

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_MUTEXATTR</b>	Invalid mutex attribute handle.
<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.
<b>PTHDB_NOSYS</b>	Not implemented
<b>PTHDB_POINTER</b>	Invalid pointer

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

**pthdb\_mutex\_addr, pthdb\_mutex\_lock\_count, pthdb\_mutex\_owner, pthdb\_mutex\_pshared, pthdb\_mutex\_prioceiling, pthdb\_mutex\_protocol, pthdb\_mutex\_state or pthdb\_mutex\_type**  
**Subroutine**

## Purpose

Gets the owner's pthread, mutex's pshared value, priority ceiling, protocol, lock state, and type.

## Library

pthread debug library (`libpthdebug.a`)

## Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_mutex_addr (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_addr_t * addrp)
```

```
int pthdb_mutex_owner (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_pthread_t * ownerp)
```

```
int pthdb_mutex_lock_count (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    int * countp);
```

```
int pthdb_mutex_pshared (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_pshared_t * psharedp)
```

```
int pthdb_mutex_prioceiling (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_pshared_t * prioceilingp)
```

```
int pthdb_mutex_protocol (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_pshared_t * protocolp)
```

```
int pthdb_mutex_state (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_mutex_state_t * statep)
```

```
int pthdb_mutex_type (pthdb_session_t session,  
                    pthdb_mutex_t mutex,  
                    pthdb_mutex_type_t * typep)
```

## Description

**pthdb\_mutex\_addr** reports the address of the `pthread_mutex_t`.

**pthdb\_mutex\_lock\_count** reports the lock count of the mutex.

**pthdb\_mutex\_owner** is used to get the pthread that owns the mutex.

The **pthdb\_mutex\_pshared** function is used to get the mutex process shared value. The pshared value can be **PSH\_SHARED**, **PSH\_PRIVATE**, or **PSH\_NOTSUP**.

**pthdb\_mutex\_prioceiling** function is used to get the mutex priority ceiling value.

**pthdb\_mutex\_protocol** function is used to get the mutex protocol value. The protocol value can be **MP\_INHERIT**, **MP\_PROTECT**, **MP\_NONE**, or **MP\_NOTSUP**.

**pthdb\_mutex\_state** is used to get the value of the mutex lock state. The state can be **MS\_LOCKED**, **MS\_UNLOCKED** or **MS\_NOTSUP**.

**pthdb\_mutex\_type** is used to get the value of the mutex type. The values for the mutex type can be **MK\_NONRECURSIVE\_NP**, **MK\_RECURSIVE\_NP**, **MK\_FAST\_NP**, **MK\_ERRORCHECK**, **MK\_RECURSIVE**, **MK\_NORMAL**, or **MK\_NOTSUP**.

## Parameters

<i>addr</i>	Mutex address
<i>countp</i>	Mutex lock count
<i>mutex</i>	Mutex handle
<i>ownerp</i>	Pointer to mutex owner
<i>psharedp</i>	Pointer to pshared value
<i>prioceilingp</i>	Pointer to priority ceiling value
<i>protocolp</i>	Pointer to protocol value
<i>session</i>	Session handle.
<i>statep</i>	Pointer to mutex state
<i>typep</i>	Pointer to mutex type

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_MUTEX</b>	Invalid mutex handle.
<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Call failed.
<b>PTHDB_NOSYS</b>	Not implemented
<b>PTHDB_POINTER</b>	Invalid pointer

## Related Information

The **pthdebug.h** file and the **pthread.h** file.

The **pthread.h** file.



---

## pthdb\_mutex\_waiter, pthdb\_cond\_waiter, pthdb\_rwlock\_read\_waiter or pthdb\_rwlock\_write\_waiter Subroutine

### Purpose

Gets the next waiter in the list of an object's waiters.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_mutex_waiter (pthdb_session_t session,
                       pthdb_mutex_t   mutex,
                       pthdb_pthread_t * waiter,
                       int               cmd);
int pthdb_cond_waiter (pthdb_session_t session,
                      pthdb_cond_t   cond,
                      pthdb_pthread_t * waiter,
                      int             cmd)
int *pthdb_rwlock_read_waiter (pthdb_session_t session,
                               pthdb_rwlock_t  rwlock,
                               pthdb_pthread_t * waiter,
                               int             cmd)
int *pthdb_rwlock_write_waiter (pthdb_session_t session,
                                pthdb_rwlock_t  rwlock,
                                pthdb_pthread_t * waiter,
                                int             cmd)
```

### Description

The **pthdb\_mutex\_waiter** functions get the next waiter in the list of an object's waiters.

Each list is reset to the top of the list when the **pthdb\_session\_update** function is called, or when the list function reports a **PTHDB\_INVALID\_\*** value. For example, when **pthdb\_attr** reports an attribute of **PTHDB\_INVALID\_ATTR** the list is reset to the beginning such that the next call reports the first attribute in the list, if any.

A report of **PTHDB\_INVALID\_OBJECT** represents the empty list or the end of a list, where *OBJECT* is one of these values: **PTHREAD**, **ATTR**, **MUTEX**, **MUTEXATTR**, **COND**, **CONDATTR**, **RWLOCK**, **RWLOCKATTR**, **KEY**, or **TID** as appropriate.

When **PTHDB\_LIST\_FIRST** is passed for the *cmd* parameter, the first item in the list is retrieved.

### Parameters

<i>session</i>	Session handle.
<i>mutex</i>	Mutex object.
<i>cond</i>	Condition variable object.
<i>cmd</i>	Reset to the beginning of the list.
<i>rwlock</i>	Read/Write lock object.
<i>waiter</i>	Pointer to waiter.

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_CMD</b>	Invalid command.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.
<b>PTHDB_MEMORY</b>	Not enough memory
<b>PTHDB_POINTER</b>	Invalid pointer

## Related Information

The **pthdebug.h** file.

The **pthread.h** file.

---

## pthdb\_pthread\_arg Subroutine

### Purpose

Reports the information associated with a pthread.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_pthread_arg (pthdb_session_t session,  
                     pthdb_pthread_t pthread,  
                     pthdb_addr_t * argp)
```

```
int pthdb_pthread_addr (pthdb_session_t session,  
                      pthdb_pthread_t pthread,  
                      pthdb_addr_t * addrp)
```

```
int pthdb_pthread_cancelpend (pthdb_session_t session,  
                             pthdb_pthread_t pthread,  
                             int * cancelpendp)
```

```
int pthdb_pthread_cancelstate (pthdb_session_t session,  
                              pthdb_pthread_t pthread,  
                              pthdb_cancelstate_t * cancelstatep)
```

```
int pthdb_pthread_canceltype (pthdb_session_t session,  
                              pthdb_pthread_t pthread,  
                              pthdb_canceltype_t * canceltypep)
```

```
int pthdb_pthread_detachstate (pthdb_session_t session,  
                               pthdb_pthread_t pthread,  
                               pthdb_detachstate_t * detachstatep)
```

```

int pthdb_thread_exit (pthdb_session_t session,
                      pthdb_thread_t pthread,
                      pthdb_addr_t * exitp)

int pthdb_thread_func (pthdb_session_t session,
                      pthdb_thread_t pthread,
                      pthdb_addr_t * funcp)

int pthdb_thread_ptid (pthdb_session_t session,
                      pthdb_thread_t pthread,
                      pthread_t * ptidp)

int pthdb_thread_schedparam (pthdb_session_t session,
                             pthdb_thread_t pthread,
                             struct sched_param * schedparamp);

int pthdb_thread_schedpolicy (pthdb_session_t session,
                              pthdb_thread_t pthread,
                              pthdb_schedpolicy_t * schedpolicy)

int pthdb_thread_schedpriority (pthdb_session_t session,
                                pthdb_thread_t pthread,
                                int * schedpriority)

int pthdb_thread_scope (pthdb_session_t session,
                       pthdb_thread_t pthread,
                       pthdb_scope_t * scopep)

int pthdb_thread_state (pthdb_session_t session,
                       pthdb_thread_t pthread,
                       pthdb_state_t * statep)

int pthdb_thread_suspendstate (pthdb_session_t session,
                               pthdb_thread_t pthread,
                               pthdb_suspendstate_t * suspendstatep)

int pthdb_ptid_thread (pthdb_session_t session,
                      pthread_t ptid,
                      pthdb_thread_t * pthreadp)

```

## Description

**pthdb\_thread\_arg** reports the initial argument passed to the pthread's start function.

**pthdb\_thread\_addr** reports the address of the pthread\_t.

**pthdb\_thread\_cancelpend** reports non-zero if cancellation is pending on the pthread; if not, it reports zero.

**pthdb\_thread\_cancelstate** reports whether cancellation is enabled (**PCS\_ENABLE**) or disabled (**PCS\_DISABLE**). **PCS\_NOTSUP** is reserved for unexpected results.

**pthdb\_thread\_canceltype** reports whether cancellation is deferred (**PCT\_DEFERRED**) or asynchronous (**PCT\_ASYNCHRONOUS**). **PCT\_NOTSUP** is reserved for unexpected results.

**pthdb\_thread\_detachstate** reports whether the pthread is detached (**PDS\_DETACHED**) or joinable (**PDS\_JOINABLE**). **PDS\_NOTSUP** is reserved for unexpected results.

**pthdb\_thread\_exit** reports the exit status returned by the pthread via **pthread\_exit**. This is only valid if the pthread has exited (**PST\_TERM**).

**pthdb\_thread\_func** reports the address of the pthread's start function.

**pthdb\_thread\_ptid** reports the pthread identifier (**pthread\_t**) associated with the pthread.

**pthdb\_thread\_schedparam** reports the pthread's scheduling parameters. This currently includes policy and priority.

**pthdb\_thread\_schedpolicy** reports whether the pthread's scheduling policy is other (**SP\_OTHER**), first in first out (**SP\_FIFO**), or round robin (**SP\_RR**). **SP\_NOTSUP** is reserved for unexpected results.

**pthdb\_thread\_schedpriority** reports the pthread's scheduling priority.

**pthdb\_thread\_scope** reports whether the pthread has process scope (**PS\_PROCESS**) or system scope (**PS\_SYSTEM**). **PS\_NOTSUP** is reserved for unexpected results.

**pthdb\_thread\_state** reports whether the pthread is being created (**PST\_IDLE**), currently running (**PST\_RUN**), waiting on an event (**PST\_SLEEP**), waiting on a cpu (**PST\_READY**), or waiting on a join or detach (**PST\_TERM**). **PST\_NOTSUP** is reserved for unexpected results.

**pthdb\_thread\_suspendstate** reports whether the pthread is suspended (**PSS\_SUSPENDED**) or not (**PSS\_UNSPENDED**). **PSS\_NOTSUP** is reserved for unexpected results.

**pthdb\_ptid\_thread** reports the pthread for the ptid.

## Parameters

<i>addr</i>	pthread address
<i>argp</i>	Initial argument buffer.
<i>cancelpendp</i>	Cancel pending buffer.
<i>cancelstatep</i>	Cancel state buffer.
<i>canceltypep</i>	Cancel type buffer.
<i>detachstatep</i>	Detach state buffer.
<i>exitp</i>	Exit value buffer.
<i>funcp</i>	Start function buffer.
<i>pthread</i>	pthread handle.
<i>pthreadp</i>	Pointer to pthread handle.
<i>ptid</i>	pthread identifier
<i>ptidp</i>	pthread identifier buffer.
<i>schedparamp</i>	Scheduling parameters buffer.
<i>schedpolicyp</i>	Scheduling policy buffer.
<i>schedpriorityp</i>	Scheduling priority buffer.
<i>scopep</i>	Contention scope buffer.
<i>session</i>	Session handle.
<i>statep</i>	State buffer.
<i>suspendstatep</i>	Suspend state buffer.

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**, else an error code is returned.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_BAD_PTID</b>	Invalid ptid.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_NOTSUP</b>	Not supported.
<b>PTHDB_INTERNAL</b>	Error in library.

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

## **pthdb\_pthread\_context** or **pthdb\_pthread\_setcontext** Subroutine

### Purpose

Provides access to the pthread context via the *struct context64* structure.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_pthread_context (pthdb_session_t  session,  
                          pthdb_pthread_t  pthread,  
                          pthdb_context_t * context)
```

```
int pthdb_pthread_setcontext (pthdb_session_t  session,  
                             pthdb_pthread_t  pthread,  
                             pthdb_context_t * context)
```

### Description

The pthread debug library provides access to the pthread context via the *struct context64* structure, whether the process is 32-bit or 64-bit. The debugger should be able to convert from 32-bit to 64-bit and from 64-bit for 32-bit processes. The extent to which this structure is filled in depends on the presence of the **PTHDB\_FLAG\_GPRS**, **PTHDB\_FLAG\_SPRSI** and **PTHDB\_FLAG\_FPRS** session flags. It is necessary to use the pthread debug library to access the context of a pthread without a kernel thread. The pthread debug library can also be used to access the context of a pthread with a kernel thread, but this results in a call back to the debugger, meaning that the debugger is capable of obtaining this information by itself. The debugger determines if the kernel thread is running in user mode or kernel mode and then fills in the *struct context64* appropriately. The pthread debug library does not use this information itself and is thus not sensitive to the correct implementation of the **read\_regs** and **write\_regs** call back functions.

**pthdb\_pthread\_context** reports the context of the pthread based on the settings of the session flags. Uses the **read\_regs** call back if the pthread has a kernel thread. If **read\_regs** is not defined, then it returns **PTHDB\_NOTSUP**.

**pthdb\_thread\_setcontext** sets the context of the pthread based on the settings of the session flags. Uses the **write\_data** call back if the pthread does not have a kernel thread. Use the **write\_regs** call back if the pthread has a kernel thread.

If the debugger does not define the **read\_regs** and **write\_regs** call backs and if the pthread does not have a kernel thread, then the **pthdb\_thread\_context** and **pthdb\_thread\_setcontext** functions succeed. But if a pthread does not have a kernel thread, then these functions fail and return **PTHDB\_CONTEXT**.

## Parameters

<i>session</i>	Session handle.
<i>pthread</i>	pthread handle.
<i>context</i>	Context buffer pointer.

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_CALLBACK</b>	Callback function failed.
<b>PTHDB_CONTEXT</b>	Could not determine pthread context.
<b>PTHDB_MEMORY</b>	Not enough memory
<b>PTHDB_NOTSUP</b>	<b>pthdb_thread_(set)context</b> returns <b>PTHDB_NOTSUP</b> if the <b>read_regs</b> , <b>write_data</b> or <b>write_regs</b> call backs are set to NULL.
<b>PTHDB_INTERNAL</b>	Error in library.

## Related Information

The **pthdebug.h** file.

The **pthread.h** file.

---

## pthdb\_thread\_hold, pthdb\_thread\_holdstate or pthdb\_thread\_unhold Subroutine

### Purpose

Reports and changes the hold state of the specified pthread.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_thread_holdstate (pthdb_session_t  session,  
                           pthdb_thread_t   pthread,
```

```

                                pthread_holdstate_t * holdstatep)
int pthread_hold (pthread_session_t session,
                 pthread_t pthread)
int pthread_unhold (pthread_session_t session,
                  pthread_t pthread)

```

## Description

**pthread\_holdstate** reports if a pthread is held. The possible hold states are **PHS\_HELD**, **PHS\_NOTHELD**, or **PHS\_NOTSUP**.

**pthread\_hold** prevents the specified pthread from running.

**pthread\_unhold** unholds the specified pthread. The pthread held earlier can be unheld by calling this function.

### Notes:

1. You must always use the **pthread\_hold** and **pthread\_unhold** functions, regardless of whether or not a pthread has a kernel thread.
2. These functions are only supported when the **PTHDB\_FLAG\_HOLD** is set.

## Parameters

<i>session</i>	Session handle.
<i>pthread</i>	pthread handle. The specified pthread should have an attached kernel thread id.
<i>holdstatep</i>	Pointer to the hold state

## Return Values

If successful, **pthread\_hold** returns **PTHDB\_SUCCESS**. Otherwise, it returns an error code.

## Error Codes

<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_HELD</b>	pthread is held.
<b>PTHDB_INTERNAL</b>	Error in library.

## Related Information

The **pthread\_session\_setflags** subroutine.

The **pthdebug.h** file.

The **pthread.h** file.

## pthread\_sigmask, pthread\_sigpend or pthread\_sigwait Subroutine

### Purpose

Returns the pthread signals pending, the signals blocked, the signals received, and awaited signals.

## Library

pthread debug library (**libpthdebug.a**)

## Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_pthread_sigmask (pthdb_session_t session,  
                           pthdb_pthread_t pthread,  
                           sigset_t * sigsetp)  
int pthdb_pthread_sigpend (pthdb_session_t session,  
                           pthdb_pthread_t pthread,  
                           sigset_t * sigsetp)  
int pthdb_pthread_sigwait (pthdb_session_t session,  
                           pthdb_pthread_t pthread,  
                           sigset_t * sigsetp)
```

## Description

**pthdb\_pthread\_sigmask** reports the signals that the pthread has blocked.

**pthdb\_pthread\_sigpend** reports the signals that the pthread has pending.

**pthdb\_pthread\_sigwait** reports the signals that the pthread is waiting on.

## Parameters

<i>session</i>	Session handle.
<i>pthread</i>	Pthread handle
<i>sigsetp</i>	Signal set buffer.

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Code

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.

## Related Information

The **pthdebug.h** file.

The **pthread.h** file.

---

## pthdb\_pthread\_specific Subroutine

### Purpose

Reports the value associated with a pthreads specific data key.



## Library

pthread debug library (**libpthdebug.a**)

## Syntax

```
#include <sys/pthdebug.h>
```

```
void *pthdb_thread_specific(pthdb_session_t session,  
                             pthdb_thread_t pthread,  
                             pthdb_key_t key,  
                             pthdb_addr_t * specificp)
```

## Description

Each process has active pthread specific data keys. Each active pthread specific data key is in use by one or more pthreads. Each pthread can have its own value associated with each pthread specific data key. The **pthdb\_thread\_specific** function provide access to those values.

**pthdb\_thread\_specific** reports the specific data value for the pthread and key combination.

## Parameters

<i>session</i>	The session handle.
<i>pthread</i>	The pthread handle.
<i>key</i>	The key.
<i>specificp</i>	Specific data value buffer.a

## Return Values

If successful, **pthdb\_thread\_specific** returns **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_KEY</b>	Invalid key.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.

## Related information

The **pthdebug.h** file.

The **pthread.h** file.

---

## pthdb\_thread\_tid or pthdb\_tid\_thread Subroutine

### Purpose

Gets the kernel thread associated with the pthread and the pthread associated with the kernel thread.

### Library

pthread debug library (**libpthdebug.a**)

## Syntax

```
#include <sys/pthdebug.h>

int pthdb_pthread_tid (pthdb_session_t session,
                      pthdb_pthread_t pthread,
                      tid_t * tidp)
int pthdb_tid_pthread (pthdb_session_t session,
                      tid_t tid,
                      pthdb_pthread_t * pthreadp)
```

## Description

**pthdb\_pthread\_tid** gets the kernel thread id associated with the pthread.

**pthdb\_tid\_pthread** is used to get the pthread associated with the kernel thread.

## Parameters

<i>session</i>	Session handle.
<i>pthread</i>	Pthread handle
<i>pthreadp</i>	Pointer to pthread handle
<i>tid</i>	Kernel thread id
<i>tidp</i>	Pointer to kernel thread id

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_PTHREAD</b>	Invalid pthread handle.
<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_TID</b>	Invalid <i>tid</i> .
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.
<b>PTHDB_INVALID_TID</b>	Empty list or the end of a list.

## Related Information

The **pthdebug.h** file.

The **pthread.h** file.

---

## pthdb\_rwlockattr\_addr, or pthdb\_rwlockattr\_pshared Subroutine

### Purpose

Gets the rwlock attribute pshared values.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthread_rwlockattr_addr (pthread_session_t    session,
                             pthread_rwlockattr_t rwlockattr,
                             pthread_addr_t      * addrp)

int pthread_rwlockattr_pshared (pthread_session_t    session,
                                pthread_rwlockattr_t rwlockattr,
                                pthread_pshared_t     * psharedp)
```

## Description

**pthread\_rwlockattr\_addr** reports the address of the `pthread_rwlockattr_t`.

**pthread\_rwlockattr\_pshared** is used to get the rwlock attribute process shared value. The pshared value can be **PSH\_SHARED**, **PSH\_PRIVATE**, or **PSH\_NOTSUP**.

## Parameters

<i>addr</i>	Read/Write lock attribute address.
<i>psharedp</i>	Pointer to the pshared value.
<i>rwlockattr</i>	Read/Write lock attribute handle
<i>session</i>	Session handle.

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

<b>PTHDB_BAD_RWLOCKATTR</b>	Invalid rwlock attribute handle.
<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.
<b>PTHDB_POINTER</b>	Invalid pointer

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

## **pthread\_rwlock\_addr, pthread\_rwlock\_lock\_count, pthread\_rwlock\_owner, pthread\_rwlock\_pshared or pthread\_rwlock\_state Subroutine**

### Purpose

Gets the owner, the pshared value, or the state of the read/write lock.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```

int pthdb_rwlock_addr (pthdb_session_t session,
                      pthdb_rwlock_t  rwlock,
                      pthdb_addr_t *  addrp)

int pthdb_rwlock_lock_count (pthdb_session_t session,
                             pthdb_rwlock_t  rwlock,
                             int * countp);

int pthdb_rwlock_owner (pthdb_session_t session,
                       pthdb_rwlock_t  rwlock,
                       pthdb_thread_t * ownerp
                       int               cmd)

int pthdb_rwlock_pshared (pthdb_session_t session,
                          pthdb_rwlock_t  rwlock,
                          pthdb_pshared_t * psharedp)

int pthdb_rwlock_state (pthdb_session_t session,
                       pthdb_rwlock_t  rwlock,
                       pthdb_rwlock_state_t * statep)

```

## Description

The **pthdb\_rwlock\_addr** function reports the address of the **pthdb\_rwlock\_t**.

The **pthdb\_rwlock\_lock\_count** function reports the lock count for the **rwlock**.

The **pthdb\_rwlock\_owner** function is used to get the read/write lock owner's pthread handle.

The **pthdb\_rwlock\_pshared** function is used to get the **rwlock** attribute process shared value. The **pshared** value can be **PSH\_SHARED**, **PSH\_PRIVATE**, or **PSH\_NOTSUP**.

The **pthdb\_rwlock\_state** is used to get the read/write locks state. The state can be **RWLS\_NOTSUP**, **RWLS\_WRITE**, **RWLS\_FREE**, and **RWLS\_READ**.

## Parameters

<i>addrp</i>	Read write lock address.
<i>countp</i>	Read write lock lock count.
<i>cmd</i>	<i>cmd</i> can be <b>PTHDB_LIST_FIRST</b> to get the first owner in the list of owners or <b>PTHDB_LIST_NEXT</b> to get the next owner in the list of owners. The list is empty or ended by <b>*owner == PTHDB_INVALID_PTHREAD</b> .
<i>ownerp</i>	Pointer to pthread which owns the <b>rwlock</b>
<i>psharedp</i>	Pointer to <b>pshared</b> value
<i>rwlock</i>	Read write lock handle
<i>session</i>	Session handle.
<i>statep</i>	Pointer to state value

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, an error code is returned.

## Error Codes

**PTHDB\_BAD\_SESSION** Invalid session handle.

<b>PTHDB_BAD_CMD</b>	Invalid command passed.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_INTERNAL</b>	Error in library.
<b>PTHDB_POINTER</b>	Invalid pointer

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

## pthdb\_session\_committed Subroutines

### Purpose

Facilitates examining and modifying multi-threaded application's pthread library object data.

### Library

pthread debug library (**libpthdebug.a**)

### Syntax

```
#include <sys/pthdebug.h>
```

```
int pthdb_session_committed (pthdb_session_t session,
                             char          ** name);
int pthdb_session_concurrency (pthdb_session_t session,
                               int           * concurrencyp);
int pthdb_session_destroy (pthdb_session_t session)
int pthdb_session_flags (pthdb_session_t session,
                         unsigned long long * flagsp)
int pthdb_session_init (pthdb_user_t user,
                       pthdb_exec_mode_t exec_mode,
                       unsigned long long flags,
                       pthdb_callbacks_t * callbacks,
                       pthdb_session_t * sessionp)
int pthdb_session_pthreaded (pthdb_user_t user,
                             unsigned long long flags
                             pthdb_callbacks_t * callbacks,
                             char          ** name)
int pthdb_session_continue_tid (pthdb_session_t session,
                                tid_t          * tidp,
                                int           cmd);
int pthdb_session_stop_tid (pthdb_session_t session,
                            tid_t          tid);
int pthdb_session_commit_tid (pthdb_session_t session,
                              tid_t          * tidp,
                              int           cmd);
int pthdb_session_setflags (pthdb_session_t session,
                            unsigned long long flags)
int pthdb_session_update (pthdb_session_t session)
```

### Description

To facilitate debugging multiple processes, the pthread debug library supports multiple sessions, one per process. Functions are provided to initialize, destroy, and customize the behavior of these sessions. In

addition, functions are provided to query global fields of the pthread library. All functions in the library require a session handle associated with an initialized session except **pthdb\_session\_init**, which initializes sessions, and **pthdb\_session\_threaded**, which can be called before the session has been initialized.

**pthdb\_session\_committed** reports the symbol name of a function called after the hold/unhold commit operation has completed. This symbol name can be used to set a breakpoint to notify the debugger when the hold/unhold commit has completed. The actual symbol name reported may change at any time. The function name returned is implemented in assembly with the following code:

```
ori 0,0, 0      # no-op
blr           # return to caller
```

This allows the debugger to overwrite the no-op with a trap instruction and leave it there by stepping over it. This function is only supported when the **PTHDB\_FLAG\_HOLD** flag is set.

**pthdb\_session\_concurrency** reports the concurrency level of the pthread library. The concurrency level is the M:N ratio, where N is always 1.

**pthdb\_session\_destroy** notifies the pthread debug library that the debugger or application is finished with the session. This deallocates any memory associated with the session and allows the session handle to be reused.

**pthdb\_session\_setflags** changes the flags for a session. With these flags, a debugger can customize the session. Flags consist of the following values or-ed together:

<b>PTHDB_FLAG_GPRS</b>	The general purpose registers should be included in any context read or write, whether internal to the library or via call backs to the debugger.
<b>PTHDB_FLAG_SPRS</b>	The special purpose registers should be included in any context read or write whether internal to the library or via call backs to the debugger.
<b>PTHDB_FLAG_FPRS</b>	The floating point registers should be included in any context read or write whether internal to the library or via call backs to the debugger.
<b>PTHDB_FLAG_REGS</b>	All registers should be included in any context read or write whether internal to the library or via call backs to the debugger. This is equivalent to <b>PTHDB_FLAG_GPRSIPTHDB_FLAG_GPRSIPTHDB_FLAG_GPRS</b> .
<b>PTHDB_FLAG_HOLD</b>	The debugger will be using the pthread debug library hold/unhold facilities to prevent the execution of pthreads. This flag cannot be used with <b>PTHDB_FLAG_SUSPEND</b> . This flag should be used by debuggers, only.
<b>PTHDB_FLAG_SUSPEND</b>	Applications will be using the pthread library suspend/continue facilities to prevent the execution of pthreads. This flag cannot be used with <b>PTHDB_FLAG_HOLD</b> . This flag is for introspective mode and should be used by applications, only. <b>Note:</b> <b>PTHDB_FLAG_HOLD</b> and <b>PTHDB_FLAG_SUSPEND</b> can only be passed to the <b>pthdb_session_init</b> function. Neither <b>PTHDB_FLAG_HOLD</b> nor <b>PTHDB_FLAG_SUSPEND</b> should be passed to <b>pthdb_session_init</b> when debugging a core file.

The **pthdb\_session\_flags** function gets the current flags for the session.

The **pthdb\_session\_init** function tells the pthread debug library to initialize a session associated with the unique given user handle. **pthdb\_session\_init** will assign a unique session handle and return it to the debugger. If the application's execution mode is 32 bit, then the debugger should initialize the **exec\_mode** to **PEM\_32BIT**. If the application's execution mode is 64 bit, then the debugger should initialize **mode** to **PEM\_64BIT**. The **flags** are documented above with the **pthdb\_session\_setflags** function. The **callback** parameter is a list of call back functions. (Also see the **pthdebug.h** header file.) The **pthdb\_session\_init** function calls the **symbol\_addrs** function to get the starting addresses of the symbols and initializes these symbols' starting addresses within the pthread debug library.

**pthdb\_session\_threaded** reports the symbol name of a function called after the pthread library has been initialized. This symbol name can be used to set a breakpoint to notify the debugger when to initialize a pthread debug library session and begin using the pthread debug library to examine pthread library state. The actual symbol name reported may change at any time. This function, is the only pthread debug library function that can be called before the pthread library is initialized. The function name returned is implemented in assembly with the following code:

```
ori 0,0,0      # no-op
blr           # return to caller
```

This is conveniently allows the debugger to overwrite the no-op with a trap instruction and leave it there by stepping over it.

The **pthdb\_session\_continue\_tid** function allows the debugger to obtain the list of threads that must be continued before it proceeds with single stepping a single pthread or continuing a group of pthreads. This function reports one tid at a time. If the list is empty or the end of the list has been reached, **PTHDB\_INVALID\_TID** is reported. The debugger will need to continue any pthreads with kernel threads that it wants. The debugger is responsible for parking the stop thread and continuing the stop thread. The *cmd* parameter can be either **PTHDB\_LIST\_NEXT** or **PTHDB\_LIST\_FIRST**; if **PTHDB\_LIST\_FIRST** is passed, then the internal counter will be reset and the first tid in the list will be reported.

**Note:** This function is only supported when the **PTHDB\_FLAG\_HOLD** flag is set.

The **pthdb\_session\_stop\_tid** function informs the pthread debug library, which informs the pthread library the tid of the thread that stopped the debugger.

**Note:** This function is only supported when the **PTHDB\_FLAG\_HOLD** flag is set.

**pthdb\_session\_commit\_tid** reports subsequent kernel thread identifiers which must be continued to commit the hold and unhold changes. This function reports one tid at a time. If the list is empty or the end of the list has been reached, **PTHDB\_INVALID\_TID** is reported. The *cmd* parameter can be either **PTHDB\_LIST\_NEXT** or **PTHDB\_LIST\_FIRST**, if **PTHDB\_LIST\_FIRST** is passed then the internal counter will be reset and first tid in the list will be reported.

**Note:** This function is only supported when the **PTHDB\_FLAG\_HOLD** flag is set.

**pthdb\_session\_update** tells the pthread debug library to update it's internal information concerning the state of the pthread library. This should be called each time the process stops before any other pthread debug library functions to ensure their results are reliable.

Each list is reset to the top of the list when the **pthdb\_session\_update** function is called, or when the list function reports a **PTHDB\_INVALID\_\*** value. For example, when **pthdb\_attr** reports an attribute of **PTHDB\_INVALID\_ATTR** the list is reset to the beginning such that the next call reports the first attribute in the list, if any.

A report of **PTHDB\_INVALID\_OBJECT** represents the empty list or the end of a list, where *OBJECT* is one of these values: **PTHREAD**, **ATTR**, **MUTEX**, **MUTEXATTR**, **COND**, **CONDATTR**, **RWLOCK**, **RWLOCKATTR**, **KEY**, or **TID** as appropriate.

## Parameters

<b>session</b>	Session handle.
<b>user</b>	Debugger user handle.
<b>sessionp</b>	Pointer to session handle.
<b>name</b>	Symbol name buffer.
<b>cmd</b>	Reset to the beginning of the list.
<b>concurrency</b>	Library concurrency buffer.

<b>flags</b>	Session flags.
<b>flagsp</b>	Pointer to session flags.
<b>exec_mode</b>	Debuggee execution mode: <b>PEM_32BIT</b> for 32-bit processes or <b>PEM_64BIT</b> for 64-bit processes.
<b>callbacks</b>	Call backs structure.
<b>tid</b>	Kernel thread id.
<b>tidp</b>	Kernel thread id buffer..

## Return Values

If successful, these functions return **PTHDB\_SUCCESS**. Otherwise, they return an error value.

## Error Codes

<b>PTHDB_BAD_SESSION</b>	Invalid session handle.
<b>PTHDB_BAD_VERSION</b>	Invalid pthread debug library or pthread library version.
<b>PTHDB_BAD_MODE</b>	Invalid execution mode.
<b>PTHDB_BAD_FLAGS</b>	Invalid session flags.
<b>PTHDB_BAD_CALLBACK</b>	Insufficient call back functions.
<b>PTHDB_BAD_CMD</b>	Invalid command.
<b>PTHDB_BAD_POINTER</b>	Invalid buffer pointer.
<b>PTHDB_BAD_USER</b>	Invalid user handle.
<b>PTHDB_CALLBACK</b>	Debugger call back error.
<b>PTHDB_MEMORY</b>	Not enough memory.
<b>PTHDB_NOSYS</b>	Function not implemented.
<b>PTHDB_NOT_PTHREADED</b>	pthread library not initialized.
<b>PTHDB_SYMBOL</b>	pthread library symbol not found.
<b>PTHDB_INTERNAL</b>	Error in library.

## Related Information

The `pthdebug.h` file.

The `pthread.h` file.

---

## pthread\_atfork Subroutine

### Purpose

Registers fork handlers.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <sys/types.h>
#include <unistd.h>
```

```
int pthread_atfork (prepare, parent, child)
void (*prepare)(void);
void (*parent)(void);
void (*child)(void);
```



## Description

The **pthread\_atfork** subroutine registers fork cleanup handlers. The *prepare* handler is called before the processing of the **fork** subroutine commences. The *parent* handler is called after the processing of the **fork** subroutine completes in the parent process. The *child* handler is called after the processing of the **fork** subroutine completes in the child process.

When the **fork** subroutine is called, only the calling thread is duplicated in the child process, but all synchronization variables are duplicated. The **pthread\_atfork** subroutine provides a way to prevent state inconsistencies and resulting deadlocks. The expected usage is that the *prepare* handler acquires all mutexes, and the two other handlers release them in the parent and child processes.

The *prepare* handlers are called in LIFO (Last In First Out) order; whereas the *parent* and *child* handlers are called in FIFO (first-in first-out) order. Thereafter, the order of calls to the **pthread\_atfork** subroutine is significant.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library.

## Parameters

<i>prepare</i>	Points to the pre-fork cleanup handler. If no pre-fork handling is desired, the value of this pointer should be set to <b>NULL</b> .
<i>parent</i>	Points to the parent post-fork cleanup handler. If no parent post-fork handling is desired, the value of this pointer should be set to <b>NULL</b> .
<i>child</i>	Points to the child post-fork cleanup handler. If no child post-fork handling is desired, the value of this pointer should be set to <b>NULL</b> .

## Return Values

Upon successful completion, the **pthread\_atfork** subroutine returns a value of zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_atfork** subroutine will fail if:

**ENOMEM** Insufficient table space exists to record the fork handler addresses.

The **pthread\_atfork** subroutine will not return an error code of **EINTR**.

## Related Information

The **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, **atexit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutine.

The “posix\_spawn or posix\_spawnp Subroutine” on page 1149.

Process Duplication and Termination in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_destroy Subroutine

### Purpose

Deletes a thread attributes object.

## Library

Threads Library (**libpthreads.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_attr_destroy (attr)
pthread_attr_t *attr;
```

## Description

The **pthread\_attr\_destroy** subroutine destroys the thread attributes object *attr*, reclaiming its storage space. It has no effect on the threads previously created with that object.

## Parameters

*attr* Specifies the thread attributes object to delete.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The **pthread\_attr\_destroy** subroutine is unsuccessful if the following is true:

**EINVAL** The *attr* parameter is not valid.

This function will not return an error code of [EINTR].

## Related Information

The **pthread\_attr\_init** (“pthread\_attr\_init Subroutine” on page 1217) subroutine, **pthread\_create** (“pthread\_create Subroutine” on page 1243) subroutine, the **pthread.h** file.

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_getguardsize or pthread\_attr\_setguardsize Subroutines

### Purpose

Gets or sets the thread guardsize attribute.

### Library

Threads Library (**libthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_attr_getguardsize (attr, guardsize)
const pthread_attr_t *attr;
size_t *guardsize;
```

```
int pthread_attr_setguardsize (attr, guardsize)
pthread_attr_t *attr;
size_t guardsize;
```

## Description

The *guardsize* attribute controls the size of the guard area for the created thread's stack. The *guardsize* attribute provides protection against overflow of the stack pointer. If a thread's stack is created with guard protection, the implementation allocates extra memory at the overflow end of the stack as a buffer against stack overflow of the stack pointer. If an application overflows into this buffer an error results (possibly in a SIGSEGV signal being delivered to the thread).

The guardsize attribute is provided to the application for two reasons:

- Overflow protection can potentially result in wasted system resources. An application that creates a large number of threads, and which knows its threads will never overflow their stack, can save system resources by turning off guard areas.
- When threads allocate large data structures on the stack, large guard areas may be needed to detect stack overflow.

The **pthread\_attr\_getguardsize** function gets the guardsize attribute in the *attr* object. This attribute is returned in the *guardsize* parameter.

The **pthread\_attr\_setguardsize** function sets the guardsize attribute in the *attr* object. The new value of this attribute is obtained from the *guardsize* parameter. If *guardsize* is zero, a guard area will not be provided for threads created with *attr*. If *guardsize* is greater than zero, a guard area of at least size *guardsize* bytes is provided for each thread created with *attr*.

A conforming implementation is permitted to round up the value contained in *guardsize* to a multiple of the configurable system variable PAGESIZE (see **sys/mman.h**). If an implementation rounds up the value of *guardsize* to a multiple of PAGESIZE, a call to **pthread\_attr\_getguardsize** specifying *attr* will store in the *guardsize* parameter the guard size specified by the previous **pthread\_attr\_setguardsize** function call. The default value of the guardsize attribute is PAGESIZE bytes. The actual value of PAGESIZE is implementation-dependent and may not be the same on all implementations.

If the *stackaddr* attribute has been set (that is, the caller is allocating and managing its own thread stacks), the *guardsize* attribute is ignored and no protection will be provided by the implementation. It is the responsibility of the application to manage stack overflow along with stack allocation and management in this case.

## Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>guardsize</i>	Controls the size of the guard area for the created thread's stack, and protects against overflow of the stack pointer.

## Return Values

If successful, the **pthread\_attr\_getguardsize** and **pthread\_attr\_setguardsize** functions return zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_attr\_getguardsize** and **pthread\_attr\_setguardsize** functions will fail if:

<b>EINVAL</b>	The attribute <i>attr</i> is invalid.
<b>EINVAL</b>	The <i>guardsize</i> parameter is invalid.

**EINVAL** The *guardsize* parameter contains an invalid value.

---

## pthread\_attr\_getinheritsched, pthread\_attr\_setinheritsched Subroutine

### Purpose

Gets and sets the **inheritsched** attribute (REALTIME THREADS).

### Syntax

```
#include <pthread.h>
#include <time.h>

int pthread_attr_getinheritsched(const pthread_attr_t *restrict attr,
                                int *restrict inheritsched);
int pthread_attr_setinheritsched(pthread_attr_t *attr,
                                int inheritsched);
```

### Description

The **pthread\_attr\_getinheritsched()** and **pthread\_attr\_setinheritsched()** functions, respectively, get and set the **inheritsched** attribute in the *attr* argument.

When the attributes objects are used by **pthread\_create()**, the **inheritsched** attribute determines how the other scheduling attributes of the created thread are set.

PTHREAD_INHERIT_SCHED	Specifies that the thread scheduling attributes is inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are ignored.
PTHREAD_EXPLICIT_SCHED	Specifies that the thread scheduling attributes are set to the corresponding values from this attributes object.

The PTHREAD\_INHERIT\_SCHED and PTHREAD\_EXPLICIT\_SCHED symbols are defined in the **<pthread.h>** header.

The following thread scheduling attributes defined by IEEE Std 1003.1-2001 are affected by the **inheritsched** attribute: scheduling policy (**schedpolicy**), scheduling parameters (**schedparam**), and scheduling contention scope (**contentionscope**).

### Application Usage

After these attributes have been set, a thread can be created with the specified attributes using **pthread\_create()**. Using these routines does not affect the current running thread.

### Return Values

If successful, the **pthread\_attr\_getinheritsched()** and **pthread\_attr\_setinheritsched()** functions return 0; otherwise, an error number is returned to indicate the error.

### Error Codes

The **pthread\_attr\_setschedpolicy()** function might fail if:

<b>EINVAL</b>	The value of <i>inheritsched</i> is not valid.
<b>ENOTSUP</b>	An attempt was made to set the attribute to an unsupported value.

These functions do not return an error code of **EINTR**.

## Related Information

“pthread\_attr\_destroy Subroutine” on page 1209, “pthread\_attr\_getscope and pthread\_attr\_setscope Subroutines” on page 1219, “pthread\_attr\_getschedparam Subroutine,” “pthread\_attr\_getschedpolicy, pthread\_attr\_setschedpolicy Subroutine” on page 1214, “pthread\_create Subroutine” on page 1243.

The **pthread.h** and **sched.h** files in *AIX 5L Version 5.3 Files Reference*.

---

## pthread\_attr\_getschedparam Subroutine

### Purpose

Returns the value of the schedparam attribute of a thread attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
#include <sys/sched.h>

int pthread_attr_getschedparam (attr, schedparam)
const pthread_attr_t *attr;
struct sched_param *schedparam;
```

### Description

The **pthread\_attr\_getschedparam** subroutine returns the value of the schedparam attribute of the thread attributes object *attr*. The schedparam attribute specifies the scheduling parameters of a thread created with this attributes object. The sched\_priority field of the **sched\_param** structure contains the priority of the thread. It is an integer value.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the cc\_r compiler used. In this case, the flag is automatically set.

### Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>schedparam</i>	Points to where the schedparam attribute value will be stored.

### Return Values

Upon successful completion, the value of the schedparam attribute is returned via the *schedparam* parameter, and 0 is returned. Otherwise, an error code is returned.

### Error Codes

The **pthread\_attr\_getschedparam** subroutine is unsuccessful if the following is true:

**EINVAL**     The *attr* parameter is not valid.

This function does not return EINTR.

## Related Information

The `pthread_attr_setschedparam` (“`pthread_attr_setschedparam Subroutine`” on page 1222) subroutine, `pthread_attr_init` (“`pthread_attr_init Subroutine`” on page 1217) subroutine, `pthread_getschedparam` (“`pthread_getschedparam Subroutine`” on page 1255) subroutine, the `pthread.h` file.

Threads Scheduling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_getschedpolicy, pthread\_attr\_setschedpolicy Subroutine

### Purpose

Gets and sets the `schedpolicy` attribute (REALTIME THREADS).

### Syntax

```
#include <pthread.h>
#include <time.h>

int pthread_attr_getschedpolicy(const pthread_attr_t *restrict attr,
                               int *restrict policy);
int pthread_attr_setschedpolicy(pthread_attr_t *attr, int policy);
```

### Description

The `pthread_attr_getschedpolicy()` and `pthread_attr_setschedpolicy()` functions, respectively, get and set the `schedpolicy` attribute in the `attr` argument.

The supported values of `policy` include `SCHED_FIFO`, `SCHED_RR`, and `SCHED_OTHER`, which are defined in the `<sched.h>` header. When threads executing with the scheduling policy `SCHED_FIFO`, `SCHED_RR`, or `SCHED_SPORADIC` are waiting on a mutex, they acquire the mutex in priority order when the mutex is unlocked.

### Application Usage

After these attributes have been set, a thread can be created with the specified attributes using `pthread_create()`. Using these routines does not affect the current running thread.

### Return Values

If successful, the `pthread_attr_getschedpolicy()` and `pthread_attr_setschedpolicy()` functions return 0; otherwise, an error number is returned to indicate the error.

### Error Codes

The `pthread_attr_setschedpolicy()` function might fail if:

<b>EINVAL</b>	The value of <code>policy</code> is not valid.
<b>ENOTSUP</b>	An attempt was made to set the attribute to an unsupported value.

These functions do not return an error code of **EINTR**.

## Related Information

“pthread\_attr\_destroy Subroutine” on page 1209, “pthread\_attr\_getscope and pthread\_attr\_setscope Subroutines” on page 1219, “pthread\_attr\_getinheritsched, pthread\_attr\_setinheritsched Subroutine” on page 1212, “pthread\_attr\_getschedparam Subroutine” on page 1213, “pthread\_create Subroutine” on page 1243.

The **pthread.h** and **time.h** files in *AIX 5L Version 5.3 Files Reference*.

---

## pthread\_attr\_getstackaddr Subroutine

### Purpose

Returns the value of the stackaddr attribute of a thread attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>

int pthread_attr_getstackaddr (attr, stackaddr)
const pthread_attr_t *attr;
void **stackaddr;
```

### Description

The **pthread\_attr\_getstackaddr** subroutine returns the value of the stackaddr attribute of the thread attributes object *attr*. This attribute specifies the stack address of the thread created with this attributes object.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the cc\_r compiler used. In this case, the flag is automatically set.

### Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>stackaddr</i>	Points to where the stackaddr attribute value will be stored.

### Return Values

Upon successful completion, the value of the stackaddr attribute is returned via the *stackaddr* parameter, and 0 is returned. Otherwise, an error code is returned.

### Error Codes

The **pthread\_attr\_getstackaddr** subroutine is unsuccessful if the following is true:

**EINVAL**     The *attr* parameter is not valid.

This function will not return EINTR.

## Related Information

The `pthread_attr_setstackaddr` (“pthread\_attr\_setstackaddr Subroutine” on page 1223) subroutine, `pthread_attr_init` (“pthread\_attr\_init Subroutine” on page 1217) subroutine, the `pthread.h` file.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_getstacksize Subroutine

### Purpose

Returns the value of the stacksize attribute of a thread attributes object.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>

int pthread_attr_getstacksize (attr, stacksize)
const pthread_attr_t *attr;
size_t *stacksize;
```

### Description

The `pthread_attr_getstacksize` subroutine returns the value of the stacksize attribute of the thread attributes object `attr`. This attribute specifies the minimum stacksize of a thread created with this attributes object. The value is given in bytes. For 32-bit compiled applications, the default stacksize is 96 KB (defined in the `pthread.h` file). For 64-bit compiled applications, the default stacksize is 192 KB (defined in the `pthread.h` file).

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

### Parameters

<code>attr</code>	Specifies the thread attributes object.
<code>stacksize</code>	Points to where the stacksize attribute value will be stored.

### Return Values

Upon successful completion, the value of the stacksize attribute is returned via the `stacksize` parameter, and 0 is returned. Otherwise, an error code is returned.

### Error Codes

The `pthread_attr_getstacksize` subroutine is unsuccessful if the following is true:

**EINVAL** The `attr` or `stacksize` parameters are not valid.

This function will not return an error code of [EINTR].



## Related Information

The `pthread_attr_setstacksize` (“pthread\_attr\_setstacksize Subroutine” on page 1224) subroutine, `pthread_attr_init` (“pthread\_attr\_init Subroutine”) subroutine, the `pthread.h` file.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_init Subroutine

### Purpose

Creates a thread attributes object and initializes it with default values.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_attr_init ( attr )  
pthread_attr_t *attr;
```

### Description

The `pthread_attr_init` subroutine creates a new thread attributes object `attr`. The new thread attributes object is initialized with the following default values:

*Always initialized*

Attribute	Default value
Detachstate	<code>PTHREAD_CREATE_JOINABLE</code>
Contention-scope	<code>PTHREAD_SCOPE_PROCESS</code> the default ensures compatibility with implementations that do not support this POSIX option.
Inheritsched	<code>PTHREAD_INHERITSCHED</code>
Schedparam	A <code>sched_param</code> structure which <code>sched_prio</code> field is set to 1, the least favored priority.
Schedpolicy	<code>SCHED_OTHER</code>
Stacksize	<code>PTHREAD_STACK_MIN</code>
Guardsize	<code>PAGESIZE</code>

The resulting attribute object (possibly modified by setting individual attribute values), when used by `pthread_create`, defines the attributes of the thread created. A single attributes object can be used in multiple simultaneous calls to `pthread_create`.

### Parameters

`attr` Specifies the thread attributes object to be created.

## Return Values

Upon successful completion, the new thread attributes object is filled with default values and returned via the *attr* parameter, and 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_attr_init` subroutine is unsuccessful if the following is true:

- EINVAL** The *attr* parameter is not valid.
- ENOMEM** There is not sufficient memory to create the thread attribute object.

This function will not return an error code of [EINTR].

## Related Information

The `pthread_attr_setdetachstate` (“`pthread_attr_getdetachstate` or `pthread_attr_setdetachstate` Subroutines”) subroutine, `pthread_attr_setstackaddr` (“`pthread_attr_setstackaddr` Subroutine” on page 1223) subroutine, `pthread_attr_setstacksize` (“`pthread_attr_setstacksize` Subroutine” on page 1224) subroutine, `pthread_create` (“`pthread_create` Subroutine” on page 1243) subroutine, `pthread_attr_destroy` (“`pthread_attr_destroy` Subroutine” on page 1209) subroutine, `pthread_attr_setguardsize` (“`pthread_attr_getguardsize` or `pthread_attr_setguardsize` Subroutines” on page 1210) subroutine.

The `pthread.h` file.

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_getdetachstate or pthread\_attr\_setdetachstate Subroutines

### Purpose

Sets and returns the value of the detachstate attribute of a thread attributes object.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_attr_setdetachstate (attr, detachstate)
pthread_attr_t *attr;
int detachstate;
```

```
int pthread_attr_getdetachstate (attr, detachstate)
const pthread_attr_t *attr;
int *detachstate;
```

## Description

The `detachstate` attribute controls whether the thread is created in a detached state. If the thread is created detached, then use of the ID of the newly created thread by the `pthread_detach` or `pthread_join` function is an error.

The `pthread_attr_setdetachstate` and `pthread_attr_getdetachstate`, respectively, set and get the `detachstate` attribute in the `attr` object.

The `detachstate` attribute can be set to either `PTHREAD_CREATE_DETACHED` or `PTHREAD_CREATE_JOINABLE`. A value of `PTHREAD_CREATE_DETACHED` causes all threads created with `attr` to be in the detached state, whereas using a value of `PTHREAD_CREATE_JOINABLE` causes all threads created with `attr` to be in the joinable state. The default value of the `detachstate` attribute is `PTHREAD_CREATE_JOINABLE`.

## Parameters

<code>attr</code>	Specifies the thread attributes object.
<code>detachstate</code>	Points to where the <code>detachstate</code> attribute value will be stored.

## Return Values

Upon successful completion, `pthread_attr_setdetachstate` and `pthread_attr_getdetachstate` return a value of `0`. Otherwise, an error number is returned to indicate the error.

The `pthread_attr_getdetachstate` function stores the value of the `detachstate` attribute in the `detachstate` parameter if successful.

## Error Codes

The `pthread_attr_setdetachstate` function will fail if:

**EINVAL** The value of `detachstate` was not valid.

The `pthread_attr_getdetachstate` and `pthread_attr_setdetachstate` functions will fail if:

**EINVAL** The attribute parameter is invalid.

These functions will not return an error code of `EINTR`.

## Related Information

The “`pthread_attr_setstackaddr` Subroutine” on page 1223, “`pthread_attr_setstacksize` Subroutine” on page 1224, “`pthread_create` Subroutine” on page 1243, and “`pthread_attr_init` Subroutine” on page 1217.

The `pthread.h` file in *AIX 5L Version 5.3 Files Reference*

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## `pthread_attr_getscope` and `pthread_attr_setscope` Subroutines

### Purpose

Gets and sets the scope attribute in the `attr` object.

## Library

Threads Library (**libpthreads.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_attr_setscope (attr, contentionscope)
pthread_attr_t *attr;
int contentionscope;
```

```
int pthread_attr_getscope (attr, contentionscope)
const pthread_attr_t *attr;
int *contentionscope;
```

## Description

The scope attribute controls whether a thread is created in system or process scope.

The **pthread\_attr\_getscope** and **pthread\_attr\_setscope** subroutines get and set the scope attribute in the **attr** object.

The scope can be set to **PTHREAD\_SCOPE\_SYSTEM** or **PTHREAD\_SCOPE\_PROCESS**. A value of **PTHREAD\_SCOPE\_SYSTEM** causes all threads created with the *attr* parameter to be in system scope, whereas a value of **PTHREAD\_SCOPE\_PROCESS** causes all threads created with the *attr* parameter to be in process scope.

The default value of the *contentionscope* parameter is **PTHREAD\_SCOPE\_PROCESS**.

## Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>contentionscope</i>	Points to where the scope attribute value will be stored.

## Return Values

Upon successful completion, the **pthread\_attr\_getscope** and **pthread\_attr\_setscope** subroutines return a value of 0. Otherwise, an error number is returned to indicate the error.

## Error Codes

<b>EINVAL</b>	The value of the attribute being set/read is not valid.
<b>ENOTSUP</b>	An attempt was made to set the attribute to an unsupported value.

## Related Information

The “**pthread\_create** Subroutine” on page 1243, and “**pthread\_attr\_init** Subroutine” on page 1217.

The **pthread.h** file in *AIX 5L Version 5.3 Files Reference*

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_getukeyset\_np or pthread\_attr\_setukeyset\_np Subroutine

### Purpose

Gets and sets the value of the active user-key-set attribute of a thread attributes object.

### Library

Threads library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
#include <sys/ukeys.h>
```

```
int pthread_attr_getukeyset_np (attr, ukeyset)
const pthread_attr_t * attr;
ukeyset_t * ukeyset;
```

### Description

The *ukeyset* parameter specifies the active user-key-set for a thread created with this attributes object. By default, newly-created threads can only access (both read and write) memory pages that have been assigned the default user-key **UKEY\_PUBLIC**. User-key-sets are not inherited across the **pthread\_create** subroutine.

The **pthread\_attr\_getukeyset\_np** subroutine gets the user-key-set attribute, while the **pthread\_attr\_setukeyset\_np** subroutine sets the user-key-set attribute in the thread attributes object specified by the *attr* parameter.

Both the **pthread\_attr\_getukeyset\_np** and the **pthread\_attr\_setukeyset\_np** subroutines will fail unless the **ukey\_enable** subroutine has been previously successfully run by a thread in the process. Refer to the Storage Protect Keys article for more details.

### Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>ukeyset</i>	Points to a location where the user-key-set attribute value is stored.

### Return Values

The **pthread\_attr\_getukeyset\_np** and **pthread\_attr\_setukeyset\_np** subroutines return a value of 0 on success. Otherwise, an error code is returned.

### Errors Codes

The **pthread\_attr\_getukeyset\_np** and **pthread\_attr\_setukeyset\_np** subroutines are unsuccessful if the following are true:

<b>EINVAL</b>	The attribute object specified by the <i>attr</i> parameter is invalid or the address pointed to by the <i>ukeyset</i> parameter is not aligned to hold a user-key-set.
<b>ENOSYS</b>	Process is not a user-key-enabled process.

In addition, the **pthread\_attr\_setukeyset\_np** subroutine is unsuccessful if the following is true:

<b>EINVAL</b>	The user-key-set value specified by the <i>ukeyset</i> parameter is not valid.
---------------	--

These functions will not return an error code of **EINTR**.

## Related Information

The **ukey\_enable** subroutine.

The **ukeyset\_init** subroutine.

The **ukeyset\_add\_key**, **ukeyset\_remove\_key**, **ukeyset\_add\_set**, **ukeyset\_remove\_set** subroutine.

The **ukeyset\_activate** subroutine.

The **ukeyset\_ismember** subroutine.

The **pthread\_attr\_init** subroutine.

The **pthread\_create** subroutine.

The **pthread\_attr\_destroy** subroutine.

---

## pthread\_attr\_setschedparam Subroutine

### Purpose

Sets the value of the schedparam attribute of a thread attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
#include <sys/sched.h>
```

```
int pthread_attr_setschedparam (attr, schedparam)
pthread_attr_t *attr;
const struct sched_param *schedparam;
```

### Description

The **pthread\_attr\_setschedparam** subroutine sets the value of the schedparam attribute of the thread attributes object *attr*. The schedparam attribute specifies the scheduling parameters of a thread created with this attributes object. The sched\_priority field of the **sched\_param** structure contains the priority of the thread.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the cc\_r compiler used. In this case, the flag is automatically set.

### Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>schedparam</i>	Points to where the scheduling parameters to set are stored. The sched_priority field must be in the range from 1 to 127, where 1 is the least favored priority, and 127 the most favored.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_attr_setschedparam` subroutine is unsuccessful if the following is true:

<b>EINVAL</b>	The <i>attr</i> parameter is not valid.
<b>ENOSYS</b>	The priority scheduling POSIX option is not implemented.
<b>ENOTSUP</b>	The value of the schedparam attribute is not supported.

## Related Information

The `pthread_attr_getschedparam` (“`pthread_attr_getschedparam Subroutine`” on page 1213) subroutine, `pthread_attr_init` (“`pthread_attr_init Subroutine`” on page 1217) subroutine, `pthread_create` (“`pthread_create Subroutine`” on page 1243) subroutine, the `pthread.h` file.

Threads Scheduling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_setstackaddr Subroutine

### Purpose

Sets the value of the `stackaddr` attribute of a thread attributes object.

### Library

Threads Library (`libpthread.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_attr_setstackaddr (attr, stackaddr)
pthread_attr_t *attr;
void *stackaddr;
```

### Description

The `pthread_attr_setstackaddr` subroutine sets the value of the `stackaddr` attribute of the thread attributes object *attr*. This attribute specifies the stack address of a thread created with this attributes object.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

A Provision has been made in `libpthread` to create guardpages for the user stack internally. This is used for debugging purposes only. By default, it is turned off and can be invoked by exporting the following environment variable:

```
AIXTHREAD_GUARDPAGES_FOR_USER_STACK=n (Where n is the decimal number of guard pages.)
```

**Note:** Even if it is exported, guard pages will only be constructed if both the `stackaddr` and `stacksize` attributes have been set by the caller for the thread. Also, the guard pages and alignment pages

will be created out of the user's stack (which will reduce the stack size). If the new stack size after creating guard pages is less than the minimum stack size (PTHREAD\_STACK\_MIN), then the guard pages will not be constructed.

## Parameters

*attr* Specifies the thread attributes object.  
*stackaddr* Specifies the stack address to set. It is a void pointer. The address that needs to be passed is not the beginning of the malloc generated address but the beginning of the stack. For example:

```
stackaddr = malloc(stacksize);  
pthread_attr_setstackaddr(&thread, stackaddr + stacksize);
```

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_attr_setstackaddr` subroutine is unsuccessful if the following is true:

**EINVAL** The *attr* parameter is not valid.  
**ENOSYS** The stack address POSIX option is not implemented.

## Related Information

The “`pthread_attr_getstackaddr` Subroutine” on page 1215, “`pthread_attr_init` Subroutine” on page 1217, `pthread.h` file.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_setstacksize Subroutine

### Purpose

Sets the value of the stacksize attribute of a thread attributes object.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>  
  
int pthread_attr_setstacksize (attr, stacksize)  
pthread_attr_t *attr;  
size_t stacksize;
```

### Description

The `pthread_attr_setstacksize` subroutine sets the value of the stacksize attribute of the thread attributes object *attr*. This attribute specifies the minimum stack size, in bytes, of a thread created with this attributes object.



The allocated stack size is always a multiple of 8K bytes, greater or equal to the required minimum stack size of 56K bytes (**PTHREAD\_STACK\_MIN**). The following formula is used to calculate the allocated stack size: if the required stack size is lower than 56K bytes, the allocated stack size is 56K bytes; otherwise, if the required stack size belongs to the range from  $(56 + (n - 1) * 16)$  K bytes to  $(56 + n * 16)$  K bytes, the allocated stack size is  $(56 + n * 16)$  K bytes.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

## Parameters

*attr* Specifies the thread attributes object.  
*stacksize* Specifies the minimum stack size, in bytes, to set. The default stack size is **PTHREAD\_STACK\_MIN**. The minimum stack size should be greater or equal than this value.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The **pthread\_attr\_setstacksize** subroutine is unsuccessful if the following is true:

**EINVAL** The *attr* parameter is not valid, or the value of the *stacksize* parameter exceeds a system imposed limit.  
**ENOSYS** The stack size POSIX option is not implemented.

## Related Information

The **pthread\_attr\_getstacksize** (“pthread\_attr\_getstacksize Subroutine” on page 1216) subroutine, **pthread\_attr\_init** (“pthread\_attr\_init Subroutine” on page 1217) subroutine, **pthread\_create** (“pthread\_create Subroutine” on page 1243) subroutine, the **pthread.h** file.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_attr\_setsuspendstate\_np and pthread\_attr\_getsuspendstate\_np Subroutine

### Purpose

Controls whether a thread is created in a suspended state.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>

int pthread_attr_setsuspendstate_np (attr, suspendstate)
pthread_attr_t *attr;
int suspendstate;
```

```
int pthread_attr_getsuspendstate_np (attr, suspendstate)
pthread_attr_t *attr;
int *suspendstate;
```

## Description

The *suspendstate* attribute controls whether the thread is created in a suspended state. If the thread is created suspended, the thread start routine will not execute until **pthread\_continue\_np** is run on the thread. The **pthread\_attr\_setsuspendstate\_np** and **pthread\_attr\_getsuspendstate\_np** routines, respectively, set and get the *suspendstate* attribute in the *attr* object.

The *suspendstate* attribute can be set to either **PTHREAD\_CREATE\_SUSPENDED\_NP** or **PTHREAD\_CREATE\_UNSUSPENDED\_NP**. A value of **PTHREAD\_CREATE\_SUSPENDED\_NP** causes all threads created with *attr* to be in the suspended state, whereas using a value of **PTHREAD\_CREATE\_UNSUSPENDED\_NP** causes all threads created with *attr* to be in the unsuspended state. The default value of the *suspendstate* attribute is **PTHREAD\_CREATE\_UNSUSPENDED\_NP**.

## Parameters

<i>attr</i>	Specifies the thread attributes object.
<i>suspendstate</i>	Points to where the <i>suspendstate</i> attribute value will be stored.

## Return Values

Upon successful completion, **pthread\_attr\_setsuspendstate\_np** and **pthread\_attr\_getsuspendstate\_np** return a value of 0. Otherwise, an error number is returned to indicate the error.

The **pthread\_attr\_getsuspendstate\_np** function stores the value of the *suspendstate* attribute in *suspendstate* if successful.

## Error Codes

The **pthread\_attr\_setsuspendstate\_np** function will fail if:

**EINVAL** The value of *suspendstate* is not valid.

---

## pthread\_barrier\_destroy or pthread\_barrier\_init Subroutine

### Purpose

Destroys or initializes a barrier object.

### Syntax

```
#include <pthread.h>

int pthread_barrier_destroy(pthread_barrier_t *barrier);
int pthread_barrier_init(pthread_barrier_t *restrict barrier,
    const pthread_barrierattr_t *restrict attr, unsigned count);
```

### Description

The **pthread\_barrier\_destroy** subroutine destroys the barrier referenced by the *barrier* parameter and releases any resources used by the barrier. The effect of subsequent use of the barrier is undefined until the barrier is reinitialized by another call to the **pthread\_barrier\_init** subroutine. An implementation can

use this subroutine to set the *barrier* parameter to an invalid value. The results are undefined if the **pthread\_barrier\_destroy** subroutine is called when any thread is blocked on the barrier, or if this function is called with an uninitialized barrier.

The **pthread\_barrier\_init** subroutine allocates any resources required to use the barrier referenced by the *barrier* parameter and initializes the barrier with attributes referenced by the *attr* parameter. If the *attr* parameter is NULL, the default barrier attributes are used; the effect is the same as passing the address of a default barrier attributes object. The results are undefined if **pthread\_barrier\_init** subroutine is called when any thread is blocked on the barrier (that is, has not returned from the **pthread\_barrier\_wait** call). The results are undefined if a barrier is used without first being initialized. The results are undefined if the **pthread\_barrier\_init** subroutine is called specifying an already initialized barrier.

The *count* argument specifies the number of threads that must call the **pthread\_barrier\_wait** subroutine before any of them successfully return from the call. The value specified by the *count* parameter must be greater than zero.

If the **pthread\_barrier\_init** subroutine fails, the barrier is not initialized and the contents of barrier are undefined.

Only the object referenced by the *barrier* parameter can be used for performing synchronization. The result of referring to copies of that object in calls to the **pthread\_barrier\_destroy** or **pthread\_barrier\_wait** subroutine is undefined.

## Return Values

Upon successful completion, these functions shall return zero; otherwise, an error number shall be returned to indicate the error.

## Error Codes

The **pthread\_barrier\_destroy** subroutine can fail if:

- |               |   |
|---------------|---|
| <b>EBUSY</b>  | The implementation has detected an attempt to destroy a barrier while it is in use (for example, while being used in a <b>pthread_barrier_wait</b> call) by another thread. |
| <b>EINVAL</b> | The value specified by barrier is invalid.  |

The **pthread\_barrier\_init**() function will fail if:

- |               |   |
|---------------|---|
| <b>EAGAIN</b> | The system lacks the necessary resources to initialize another barrier. |
| <b>EINVAL</b> | The value specified by the <i>count</i> parameter is equal to zero.     |
| <b>ENOMEM</b> | Insufficient memory exists to initialize the barrier.                   |

The **pthread\_barrier\_init** subroutine can fail if:

- |               |  |
|---------------|--|
| <b>EBUSY</b>  | The implementation has detected an attempt to reinitialize a barrier while it is in use (for example, while being used in a <b>pthread_barrier_wait</b> call) by another thread. |
| <b>EINVAL</b> | The value specified by the <i>attr</i> parameter is invalid.   |

## Related Information

The “pthread\_barrier\_wait Subroutine” on page 1228, “pthread\_barrierattr\_destroy or pthread\_barrierattr\_init Subroutine” on page 1229, “pthread\_barrierattr\_getpshared or pthread\_barrierattr\_setpshared Subroutine” on page 1230.

The **pthread.h** file.

---

## pthread\_barrier\_wait Subroutine

### Purpose

Synchronizes threads at a barrier.

### Syntax

```
#include <pthread.h>

int pthread_barrier_wait(pthread_barrier_t *barrier);
```

### Description

The **pthread\_barrier\_wait** subroutine synchronizes participating threads at the barrier referenced by **barrier**. The calling thread blocks until the required number of threads have called **pthread\_barrier\_wait** specifying the barrier.

When the required number of threads have called **pthread\_barrier\_wait** specifying the barrier, the constant **PTHREAD\_BARRIER\_SERIAL\_THREAD** is returned to one unspecified thread and 0 is returned to the remaining threads. At this point, the barrier resets to the state it had as a result of the most recent **pthread\_barrier\_init** function that referenced it.

The constant **PTHREAD\_BARRIER\_SERIAL\_THREAD** is defined in **<pthread.h>**, and its value is distinct from any other value returned by **pthread\_barrier\_wait**.

The results are undefined if this function is called with an uninitialized barrier.

If a signal is delivered to a thread blocked on a barrier, upon return from the signal handler, the thread resumes waiting at the barrier if the barrier wait has not completed (that is, if the required number of threads have not arrived at the barrier during the execution of the signal handler); otherwise, the thread continues as normal from the completed barrier wait. Until the thread in the signal handler returns from it, other threads might proceed past the barrier after they have all reached it.

**Note:** In AIX 5.3, when the required number of threads has called **pthread\_barrier\_wait**, the **PTHREAD\_BARRIER\_SERIAL\_THREAD** constant is returned by the last pthread that called **pthread\_barrier\_wait**. Furthermore, if a thread is in a signal handler while waiting and all the required threads have reached the barrier, the other threads can proceed past the barrier.

A thread that has blocked on a barrier does not prevent any unblocked thread that is eligible to use the same processing resources from eventually making forward progress in its execution. Eligibility for processing resources is determined by the scheduling policy.

### Parameters

*barrier* Points to the barrier where participating threads wait.

### Return Values

Upon successful completion, **pthread\_barrier\_wait** returns **PTHREAD\_BARRIER\_SERIAL\_THREAD** for a single (arbitrary) thread synchronized at the barrier and 0 for the other threads. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_barrier_destroy` subroutine can fail if:

**EINVAL**           The value specified by *barrier* does not refer to an initialized barrier object.

This function does not return an error code of **EINTR**.

## Related Information

The “`pthread_barrier_destroy` or `pthread_barrier_init` Subroutine” on page 1226, “`pthread_barrierattr_destroy` or `pthread_barrierattr_init` Subroutine,” “`pthread_barrierattr_getpshared` or `pthread_barrierattr_setpshared` Subroutine” on page 1230.

The `pthread.h` file.

---

## `pthread_barrierattr_destroy` or `pthread_barrierattr_init` Subroutine

### Purpose

Destroys or initializes the barrier attributes object.

### Syntax

```
#include <pthread.h>

int pthread_barrierattr_destroy(pthread_barrierattr_t *attr);
int pthread_barrierattr_init(pthread_barrierattr_t *attr);
```

### Description

The `pthread_barrierattr_destroy` subroutine destroys a barrier attributes object. A destroyed *attr* attributes object can be reinitialized using the `pthread_barrierattr_init` subroutine; the results of otherwise referencing the object after it has been destroyed are undefined. An implementation can cause the `pthread_barrierattr_destroy` subroutine to set the object referenced by the *attr* parameter to an invalid value.

The `pthread_barrierattr_init` subroutine initializes a barrier attributes object *attr* with the default value for all of the attributes defined by the implementation.

Results are undefined if the `pthread_barrierattr_init` subroutine is called specifying an already initialized *attr* attributes object.

After a barrier attributes object has been used to initialize one or more barriers, any function affecting the attributes object (including destruction) do not affect any previously initialized barrier.

### Return Values

If successful, the `pthread_barrierattr_destroy` and `pthread_barrierattr_init` subroutines return zero; otherwise, an error number shall be returned to indicate the error.

## Error Codes

The `pthread_barrierattr_destroy` subroutine can fail if:

**EINVAL**           The value specified by the *attr* parameter is invalid.

The `pthread_barrierattr_init` subroutine will fail if:

**ENOMEM** Insufficient memory exists to initialize the barrier attributes object.

## Related Information

The “`pthread_barrier_destroy` or `pthread_barrier_init` Subroutine” on page 1226, “`pthread_barrier_wait` Subroutine” on page 1228, “`pthread_barrierattr_getpshared` or `pthread_barrierattr_setpshared` Subroutine.”

---

## pthread\_barrierattr\_getpshared or pthread\_barrierattr\_setpshared Subroutine

### Purpose

Gets and sets the process-shared attribute of the barrier attributes object.

### Syntax

```
#include <pthread.h>

int pthread_barrierattr_getpshared(const pthread_barrierattr_t *
    restrict attr, int *restrict pshared);
int pthread_barrierattr_setpshared(pthread_barrierattr_t *attr,
    int pshared);
```

### Description

The `pthread_barrierattr_getpshared` subroutine obtains the value of the process-shared attribute from the attributes object referenced by the `attr` parameter. The `pthread_barrierattr_setpshared` subroutine sets the process-shared attribute in an initialized attributes object referenced by the `attr` parameter.

The process-shared attribute is set to **PTHREAD\_PROCESS\_SHARED** to permit a barrier to be operated upon by any thread that has access to the memory where the barrier is allocated. If the process-shared attribute is **PTHREAD\_PROCESS\_PRIVATE**, the barrier is only operated upon by threads created within the same process as the thread that initialized the barrier; if threads of different processes attempt to operate on such a barrier, the behavior is undefined. The default value of the attribute is **PTHREAD\_PROCESS\_PRIVATE**. Both constants **PTHREAD\_PROCESS\_SHARED** and **PTHREAD\_PROCESS\_PRIVATE** are defined in the `pthread.h` file.

Additional attributes, their default values, and the names of the associated functions to get and set those attribute values are implementation-defined.

### Return Values

If successful, the `pthread_barrierattr_getpshared` subroutine will return zero and store the value of the process-shared attribute of `attr` into the object referenced by the `pshared` parameter. Otherwise, an error number shall be returned to indicate the error.

If successful, the `pthread_barrierattr_setpshared` subroutine will return zero; otherwise, an error number shall be returned to indicate the error.

### Error Codes

These functions may fail if:

**EINVAL** The value specified by `attr` is invalid.

The `pthread_barrierattr_setpshared` subroutine will fail if:

**EINVAL**            The new value specified for the process-shared attribute is not one of the legal values `PTHREAD_PROCESS_SHARED` or `PTHREAD_PROCESS_PRIVATE`.

## Related Information

The “`pthread_barrier_destroy` or `pthread_barrier_init` Subroutine” on page 1226, “`pthread_barrier_wait` Subroutine” on page 1228, “`pthread_barrierattr_destroy` or `pthread_barrierattr_init` Subroutine” on page 1229.

---

## pthread\_cancel Subroutine

### Purpose

Requests the cancellation of a thread.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_cancel (thread)
pthread_t thread;
```

### Description

The `pthread_cancel` subroutine requests the cancellation of the thread *thread*. The action depends on the cancelability of the target thread:

- If its cancelability is disabled, the cancellation request is set pending.
- If its cancelability is deferred, the cancellation request is set pending till the thread reaches a cancellation point.
- If its cancelability is asynchronous, the cancellation request is acted upon immediately; in some cases, it may result in unexpected behavior.

The cancellation of a thread terminates it safely, using the same termination procedure as the `pthread_exit` (“`pthread_exit` Subroutine” on page 1248) subroutine.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

### Parameters

*thread*            Specifies the thread to be canceled.

### Return Values

If successful, the `pthread_cancel` function returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_cancel` function may fail if:

**ESRCH** No thread could be found corresponding to that specified by the given thread ID.

The `pthread_cancel` function will not return an error code of `EINTR`.

## Related Information

The `pthread_kill` (“`pthread_kill Subroutine`” on page 1265) subroutine, `pthread_exit` (“`pthread_exit Subroutine`” on page 1248) subroutine, `pthread_join` (“`pthread_join or pthread_detach Subroutine`” on page 1262) subroutine, `pthread_cond_wait` (“`pthread_cond_wait or pthread_cond_timedwait Subroutine`” on page 1236), and `pthread_cond_timedwait` (“`pthread_cond_wait or pthread_cond_timedwait Subroutine`” on page 1236) subroutines.

The `pthread.h` file.

Terminating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_cleanup\_pop or pthread\_cleanup\_push Subroutine

### Purpose

Activates and deactivates thread cancellation handlers.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
void pthread_cleanup_pop (execute)  
int execute;
```

```
void pthread_cleanup_push (routine, arg)  
void (*routine)(void *);  
void *arg;
```

### Description

The `pthread_cleanup_push` subroutine pushes the specified cancellation cleanup handler *routine* onto the calling thread's cancellation cleanup stack. The cancellation cleanup handler is popped from the cancellation cleanup stack and invoked with the argument *arg* when: (a) the thread exits (that is, calls `pthread_exit`), (b) the thread acts upon a cancellation request, or (c) the thread calls `pthread_cleanup_pop` with a nonzero *execute* argument.

The `pthread_cleanup_pop` subroutine removes the subroutine at the top of the calling thread's cancellation cleanup stack and optionally invokes it (if *execute* is nonzero).

These subroutines may be implemented as macros and will appear as statements and in pairs within the same lexical scope (that is, the `pthread_cleanup_push` macro may be thought to expand to a token list whose first token is `'{'` with `pthread_cleanup_pop` expanding to a token list whose last token is the corresponding `'}'`).



The effect of calling **longjmp** or **siglongjmp** is undefined if there have been any calls to **pthread\_cleanup\_push** or **pthread\_cleanup\_pop** made without the matching call since the jump buffer was filled. The effect of calling **longjmp** or **siglongjmp** from inside a cancellation cleanup handler is also undefined unless the jump buffer was also filled in the cancellation cleanup handler.

## Parameters

*execute*        Specifies if the popped subroutine will be executed.  
*routine*        Specifies the address of the cancellation subroutine.  
*arg*            Specifies the argument passed to the cancellation subroutine.

## Related Information

The **pthread\_cancel** (“pthread\_cancel Subroutine” on page 1231), **pthread\_setcancelstate** (“pthread\_setcancelstate, pthread\_setcanceltype, or pthread\_testcancel Subroutines” on page 1296) subroutines, the **pthread.h** file.

Terminating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_cond\_destroy or pthread\_cond\_init Subroutine

### Purpose

Initialize and destroys condition variables.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>

int pthread_cond_init (cond, attr)
pthread_cond_t *cond;
const pthread_condattr_t *attr;

int pthread_cond_destroy (cond)
pthread_cond_t *cond;

pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

### Description

The function **pthread\_cond\_init** initializes the condition variable referenced by *cond* with attributes referenced by *attr*. If *attr* is NULL, the default condition variable attributes are used; the effect is the same as passing the address of a default condition variable attributes object. Upon successful initialization, the state of the condition variable becomes initialized.

Attempting to initialize an already initialized condition variable results in undefined behavior.

The function **pthread\_cond\_destroy** destroys the given condition variable specified by *cond*; the object becomes, in effect, uninitialized. An implementation may cause **pthread\_cond\_destroy** to set the object referenced by *cond* to an invalid value. A destroyed condition variable object can be re-initialized using **pthread\_cond\_init**; the results of otherwise referencing the object after it has been destroyed are undefined.

It is safe to destroy an initialized condition variable upon which no threads are currently blocked. Attempting to destroy a condition variable upon which other threads are currently blocked results in undefined behavior.

In cases where default condition variable attributes are appropriate, the macro `PTHREAD_COND_INITIALIZER` can be used to initialize condition variables that are statically allocated. The effect is equivalent to dynamic initialization by a call to `pthread_cond_init` with parameter *attr* specified as `NULL`, except that no error checks are performed.

## Parameters

<i>cond</i>	Pointer to the condition variable.
<i>attr</i>	Specifies the attributes of the condition.

## Return Values

If successful, the `pthread_cond_init` and `pthread_cond_destroy` functions return zero. Otherwise, an error number is returned to indicate the error. The `EBUSY` and `EINVAL` error checks, if implemented, act as if they were performed immediately at the beginning of processing for the function and caused an error return prior to modifying the state of the condition variable specified by *cond*.

## Error Codes

The `pthread_cond_init` function will fail if:

<b>EAGAIN</b>	The system lacked the necessary resources (other than memory) to initialize another condition variable.
<b>ENOMEM</b>	Insufficient memory exists to initialize the condition variable.

The `pthread_cond_init` function may fail if:

<b>EINVAL</b>	The value specified by <i>attr</i> is invalid.
---------------	--

The `pthread_cond_destroy` function may fail if:

<b>EBUSY</b>	The implementation has detected an attempt to destroy the object referenced by <i>cond</i> while it is referenced (for example, while being used in a <code>pthread_cond_wait</code> or <code>pthread_cond_timedwait</code> by another thread.
<b>EINVAL</b>	The value specified by <i>cond</i> is invalid.

These functions will not return an error code of `EINTR`.

## Related Information

The `pthread_cond_signal` or `pthread_cond_broadcast` (“`pthread_cond_signal` or `pthread_cond_broadcast` Subroutine” on page 1235) subroutine and the `pthread_cond_wait` or `pthread_cond_timewait` (“`pthread_cond_wait` or `pthread_cond_timedwait` Subroutine” on page 1236) subroutine.

The `pthread.h` file.

Using Condition Variables in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## PTHREAD\_COND\_INITIALIZER Macro

### Purpose

Initializes a static condition variable with default attributes.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
static pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

### Description

The **PTHREAD\_COND\_INITIALIZER** macro initializes the static condition variable *cond*, setting its attributes to default values. This macro should only be used for static condition variables, since no error checking is performed.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

### Related Information

The **pthread\_cond\_init** (“**pthread\_cond\_destroy** or **pthread\_cond\_init** Subroutine” on page 1233) subroutine.

Using Condition Variables in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_cond\_signal or pthread\_cond\_broadcast Subroutine

### Purpose

Unblocks one or more threads blocked on a condition.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>

int pthread_cond_signal (condition)
pthread_cond_t *condition;

int pthread_cond_broadcast (condition)
pthread_cond_t *condition;
```

### Description

These subroutines unblock one or more threads blocked on the condition specified by *condition*. The **pthread\_cond\_signal** subroutine unblocks at least one blocked thread, while the **pthread\_cond\_broadcast** subroutine unblocks all the blocked threads.

If more than one thread is blocked on a condition variable, the scheduling policy determines the order in which threads are unblocked. When each thread unblocked as a result of a **pthread\_cond\_signal** or **pthread\_cond\_broadcast** returns from its call to **pthread\_cond\_wait** or **pthread\_cond\_timedwait**, the thread owns the mutex with which it called **pthread\_cond\_wait** or **pthread\_cond\_timedwait**. The thread(s) that are unblocked contend for the mutex according to the scheduling policy (if applicable), and as if each had called **pthread\_mutex\_lock**.

The **pthread\_cond\_signal** or **pthread\_cond\_broadcast** functions may be called by a thread whether or not it currently owns the mutex that threads calling **pthread\_cond\_wait** or **pthread\_cond\_timedwait** have associated with the condition variable during their waits; however, if predictable scheduling behavior is required, then that mutex is locked by the thread calling **pthread\_cond\_signal** or **pthread\_cond\_broadcast**.

If no thread is blocked on the condition, the subroutine succeeds, but the signalling of the condition is not held. The next thread calling **pthread\_cond\_wait** will be blocked.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

## Parameters

*condition*            Specifies the condition to signal.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Code

The **pthread\_cond\_signal** and **pthread\_cond\_broadcast** subroutines are unsuccessful if the following is true:

**EINVAL**        The *condition* parameter is not valid.

## Related Information

The **pthread\_cond\_wait** or **pthread\_cond\_timedwait** (“**pthread\_cond\_wait** or **pthread\_cond\_timedwait** Subroutine”) subroutine.

Using Condition Variables in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## **pthread\_cond\_wait** or **pthread\_cond\_timedwait** Subroutine

### Purpose

Blocks the calling thread on a condition.

### Library

Threads Library (**libpthreads.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_cond_wait (cond, mutex)
pthread_cond_t *cond;
pthread_mutex_t *mutex;
```

```
int pthread_cond_timedwait (cond, mutex, timeout)
pthread_cond_t *cond;
pthread_mutex_t *mutex;
const struct timespec *timeout;
```

## Description

The **pthread\_cond\_wait** and **pthread\_cond\_timedwait** functions are used to block on a condition variable. They are called with *mutex* locked by the calling thread or undefined behavior will result.

These functions atomically release *mutex* and cause the calling thread to block on the condition variable *cond*; atomically here means atomically with respect to access by another thread to the mutex and then the condition variable. That is, if another thread is able to acquire the mutex after the about-to-block thread has released it, then a subsequent call to **pthread\_cond\_signal** or **pthread\_cond\_broadcast** in that thread behaves as if it were issued after the about-to-block thread has blocked.

Upon successful return, the mutex is locked and owned by the calling thread.

When using condition variables there is always a boolean predicate involving shared variables associated with each condition wait that is true if the thread should proceed. Spurious wakeups from the **pthread\_cond\_wait** or **pthread\_cond\_timedwait** functions may occur. Since the return from **pthread\_cond\_wait** or **pthread\_cond\_timedwait** does not imply anything about the value of this predicate, the predicate should be reevaluated upon such return.

The effect of using more than one mutex for concurrent **pthread\_cond\_wait** or **pthread\_cond\_timedwait** operations on the same condition variable is undefined; that is, a condition variable becomes bound to a unique mutex when a thread waits on the condition variable, and this (dynamic) binding ends when the wait returns.

A condition wait (whether timed or not) is a cancellation point. When the cancelability enable state of a thread is set to `PTHREAD_CANCEL_DEFERRED`, a side effect of acting upon a cancellation request while in a condition wait is that the mutex is (in effect) reacquired before calling the first cancellation cleanup handler. The effect is as if the thread were unblocked, allowed to execute up to the point of returning from the call to **pthread\_cond\_wait** or **pthread\_cond\_timedwait**, but at that point notices the cancellation request and instead of returning to the caller of **pthread\_cond\_wait** or **pthread\_cond\_timedwait**, starts the thread cancellation activities, which includes calling cancellation cleanup handlers.

A thread that has been unblocked because it has been canceled while blocked in a call to **pthread\_cond\_wait** or **pthread\_cond\_timedwait** does not consume any condition signal that may be directed concurrently at the condition variable if there are other threads blocked on the condition variable.

The **pthread\_cond\_timedwait** function is the same as **pthread\_cond\_wait** except that an error is returned if the absolute time specified by *timeout* passes (that is, system time equals or exceeds *timeout*) before the condition *cond* is signaled or broadcast, or if the absolute time specified by *timeout* has already been passed at the time of the call. When such time-outs occur, **pthread\_cond\_timedwait** will nonetheless release and reacquire the mutex referenced by *mutex*. The function **pthread\_cond\_timedwait** is also a cancellation point. The absolute time specified by *timeout* can be

either based on the system realtime clock or the system monotonic clock. The reference clock for the condition variable is set by calling **pthread\_condattr\_setclock** before its initialization with the corresponding condition attributes object.

If a signal is delivered to a thread waiting for a condition variable, upon return from the signal handler the thread resumes waiting for the condition variable as if it was not interrupted, or it returns zero due to spurious wakeup.

## Parameters

<i>cond</i>	Specifies the condition variable to wait on.
<i>mutex</i>	Specifies the mutex used to protect the condition variable. The mutex must be locked when the subroutine is called.
<i>timeout</i>	Points to the absolute time structure specifying the blocked state timeout.

## Return Values

Except in the case of ETIMEDOUT, all these error checks act as if they were performed immediately at the beginning of processing for the function and cause an error return, in effect, prior to modifying the state of the mutex specified by *mutex* or the condition variable specified by *cond*.

Upon successful completion, a value of zero is returned. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_cond\_timedwait** function will fail if:

**ETIMEDOUT** The time specified by *timeout* to **pthread\_cond\_timedwait** has passed.

The **pthread\_cond\_wait** and **pthread\_cond\_timedwait** functions may fail if:

<b>EINVAL</b>	The value specified by <i>cond</i> , <i>mutex</i> , or <i>timeout</i> is invalid.
<b>EINVAL</b>	Different mutexes were supplied for concurrent <b>pthread_cond_wait</b> or <b>pthread_cond_timedwait</b> operations on the same condition variable.
<b>EINVAL</b>	The mutex was not owned by the current thread at the time of the call.
<b>EPERM</b>	The mutex was not owned by the current thread at the time of the call, XPG_SUS_ENV is set to ON, and XPG_UNIX98 is not set.

These functions will not return an error code of EINTR.

## Related Information

The **pthread\_cond\_signal** or **pthread\_cond\_broadcast** (“pthread\_cond\_signal or pthread\_cond\_broadcast Subroutine” on page 1235) subroutine, “pthread\_condattr\_getclock, pthread\_condattr\_setclock Subroutine” on page 1240, the **pthread.h** file.

Using Condition Variables in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_condattr\_destroy or pthread\_condattr\_init Subroutine

### Purpose

Initializes and destroys condition variable.

## Library

Threads Library (**libpthread.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_condattr_destroy (attr)
pthread_condattr_t *attr;
```

```
int pthread_condattr_init (attr)
pthread_condattr_t *attr;
```

## Description

The function **pthread\_condattr\_init** initializes a condition variable attributes object *attr* with the default value for all of the attributes defined by the implementation. Attempting to initialize an already initialized condition variable attributes object results in undefined behavior.

After a condition variable attributes object has been used to initialize one or more condition variables, any function affecting the attributes object (including destruction) does not affect any previously initialized condition variables.

The **pthread\_condattr\_destroy** function destroys a condition variable attributes object; the object becomes, in effect, uninitialized. The **pthread\_condattr\_destroy** subroutine may set the object referenced by *attr* to an invalid value. A destroyed condition variable attributes object can be re-initialized using **pthread\_condattr\_init**; the results of otherwise referencing the object after it has been destroyed are undefined.

## Parameter

*attr*      Specifies the condition attributes object to delete.

## Return Values

If successful, the **pthread\_condattr\_init** and **pthread\_condattr\_destroy** functions return zero. Otherwise, an error number is returned to indicate the error.

## Error Code

The **pthread\_condattr\_init** function will fail if:

**ENOMEM**    Insufficient memory exists to initialize the condition variable attributes object.

The **pthread\_condattr\_destroy** function may fail if:

**EINVAL**     The value specified by *attr* is invalid.

These functions will not return an error code of **EINTR**.

## Related Information

The **pthread\_cond\_init** (“**pthread\_cond\_destroy** or **pthread\_cond\_init** Subroutine” on page 1233) subroutine, **pthread\_condattr\_getpshared** (“**pthread\_condattr\_getpshared** Subroutine” on page 1241) subroutine, **pthread\_create** (“**pthread\_create** Subroutine” on page 1243) subroutine, **pthread\_mutex\_init** (“**pthread\_mutex\_init** or **pthread\_mutex\_destroy** Subroutine” on page 1267) subroutine.

The **pthread.h** file.

Using Condition Variables in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_condattr\_getclock, pthread\_condattr\_setclock Subroutine

### Purpose

Gets and sets the clock selection condition variable attribute.

### Syntax

```
int pthread_condattr_getclock(const pthread_condattr_t *restrict attr,  
                             clockid_t *restrict clock_id);  
int pthread_condattr_setclock(pthread_condattr_t *attr,  
                              clockid_t clock_id);
```

### Description

The **pthread\_condattr\_getclock** subroutine obtains the value of the clock attribute from the attributes object referenced by the *attr* argument. The **pthread\_condattr\_setclock** subroutine sets the clock attribute in an initialized attributes object referenced by the *attr* argument. If **pthread\_condattr\_setclock** is called with a *clock\_id* argument that refers to a CPU-time clock, the call will fail.

The clock attribute is the clock ID of the clock that shall be used to measure the timeout service of the **pthread\_cond\_timedwait** subroutine. The default value of the clock attribute refers to the system clock.

### Parameters

*attr* Specifies the condition attributes object.  
*clock\_id* For **pthread\_condattr\_getclock()**, points to where the clock attribute value will be stored.  
For **pthread\_condattr\_setclock()**, specifies the clock to set. Valid values are:

#### **CLOCK\_REALTIME**

The system realtime clock.

#### **CLOCK\_MONOTONIC**

The system monotonic clock. The value of this clock represents the amount of time since an unspecified point in the past. The value of this clock always grows: it cannot be set by **clock\_settime()** and cannot have backward clock jumps.

### Return Values

If successful, the **pthread\_condattr\_getclock** subroutine returns 0 and stores the value of the clock attribute of *attr* in the object referenced by the *clock\_id* argument. Otherwise, an error code is returned to indicate the error.

If successful, the **pthread\_condattr\_setclock** subroutine returns 0; otherwise, an error code is returned to indicate the error.

### Error Codes

**EINVAL** The value specified by *attr* is invalid.  
**EINVAL** The **pthread\_condattr\_setclock** subroutine returns this error if the value specified by the *clock\_id* does not refer to a known clock, or is a CPU-time clock.  
**ENOTSUP** The function is not supported with checkpoint-restart processes.



## Related Information

“pthread\_cond\_destroy or pthread\_cond\_init Subroutine” on page 1233, “pthread\_cond\_wait or pthread\_cond\_timedwait Subroutine” on page 1236, “pthread\_condattr\_getpshared Subroutine,” “pthread\_condattr\_destroy or pthread\_condattr\_init Subroutine” on page 1238, “pthread\_condattr\_setpshared Subroutine” on page 1242, “pthread\_create Subroutine” on page 1243, “pthread\_mutex\_init or pthread\_mutex\_destroy Subroutine” on page 1267.

The **pthread.h** file.

The Base Definitions volume of IEEE Std 1003.1-2001.

---

## pthread\_condattr\_getpshared Subroutine

### Purpose

Returns the value of the pshared attribute of a condition attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_condattr_getpshared (attr, pshared)
const pthread_condattr_t *attr;
int *pshared;
```

### Description

The **pthread\_condattr\_getpshared** subroutine returns the value of the pshared attribute of the condition attribute object *attr*. This attribute specifies the process sharing of the condition variable created with this attributes object. It may have one of the following values:

<b>PTHREAD_PROCESS_SHARED</b>	Specifies that the condition variable can be used by any thread that has access to the memory where it is allocated, even if these threads belong to different processes.
<b>PTHREAD_PROCESS_PRIVATE</b>	Specifies that the condition variable shall only be used by threads within the same process as the thread that created it. This is the default value.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

### Parameters

*attr* Specifies the condition attributes object.  
*pshared* Points to where the pshared attribute value will be stored.

### Return Values

Upon successful completion, the value of the pshared attribute is returned via the *pshared* parameter, and 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_condattr_getpshared` subroutine is unsuccessful if the following is true:

- EINVAL** The *attr* parameter is not valid.
- ENOSYS** The process sharing POSIX option is not implemented.

## Related Information

The `pthread_condattr_setpshared` (“`pthread_condattr_setpshared` Subroutine”) subroutine, `pthread_condattr_init` (“`pthread_condattr_destroy` or `pthread_condattr_init` Subroutine” on page 1238) subroutine.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_condattr\_setpshared Subroutine

### Purpose

Sets the value of the pshared attribute of a condition attributes object.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>

int pthread_condattr_setpshared (attr, pshared)
pthread_condattr_t *attr;
int pshared;
```

### Description

The `pthread_condattr_setpshared` subroutine sets the value of the pshared attribute of the condition attributes object *attr*. This attribute specifies the process sharing of the condition variable created with this attributes object.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

### Parameters

- attr* Specifies the condition attributes object.
- pshared* Specifies the process sharing to set. It must have one of the following values:

#### **PTHREAD\_PROCESS\_SHARED**

Specifies that the condition variable can be used by any thread that has access to the memory where it is allocated, even if these threads belong to different processes.

#### **PTHREAD\_PROCESS\_PRIVATE**

Specifies that the condition variable shall only be used by threads within the same process as the thread that created it. This is the default value.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_condattr_setpshared` subroutine is unsuccessful if the following is true:

**EINVAL** The `attr` or `pshared` parameters are not valid.

## Related Information

The `pthread_condattr_getpshared` (“`pthread_condattr_getpshared Subroutine`” on page 1241) subroutine, `pthread_condattr_init` or `pthread_cond_init` (“`pthread_condattr_destroy` or `pthread_condattr_init Subroutine`” on page 1238) subroutine.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_create Subroutine

### Purpose

Creates a new thread, initializes its attributes, and makes it runnable.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_create (thread, attr, start_routine (void), arg)
pthread_t *thread;
const pthread_attr_t *attr;
void **start_routine (void);
void *arg;
```

### Description

The `pthread_create` subroutine creates a new thread and initializes its attributes using the thread attributes object specified by the `attr` parameter. The new thread inherits its creating thread’s signal mask; but any pending signal of the creating thread will be cleared for the new thread.

The new thread is made runnable, and will start executing the `start_routine` routine, with the parameter specified by the `arg` parameter. The `arg` parameter is a void pointer; it can reference any kind of data. It is not recommended to cast this pointer into a scalar data type (`int` for example), because the casts may not be portable.

After thread creation, the thread attributes object can be reused to create another thread, or deleted.

The thread terminates in the following cases:

- The thread returned from its starting routine (the `main` routine for the initial thread)
- The thread called the `pthread_exit` (“`pthread_exit Subroutine`” on page 1248) subroutine

- The thread was canceled
- The thread received a signal that terminated it
- The entire process is terminated due to a call to either the **exec** (“exec: execl, execl, execlp, execv, execve, execvp, or exec Subroutine” on page 236) or **exit** (“exit, atexit, unatexit, \_exit, or \_Exit Subroutine” on page 243) subroutines.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

When multiple threads are created in a process, the **FULL\_CORE** flag is set for all signals. This means that if a core file is produced, it will be much bigger than a single\_threaded application. This is necessary to debug multiple-threaded processes.

When a process uses the **pthread\_create** function, and thus becomes multi-threaded, the **FULL\_CORE** flag is enabled for all signals. If a signal is received whose action is to terminate the process with a core dump, a full dump (usually much larger than a regular dump) will be produced. This is necessary so that multi-threaded programs can be debugged with the **dbx** command.

The following piece of pseudocode is an example of how to avoid getting a full core. Please note that in this case, debug will not be possible. It may be easier to limit the size of the core with the **ulimit** command.

```
struct sigaction siga;
siga.sa_handler = SIG_DFL;
siga.sa_flags = SA_RESTART;
SIGINITSET(siga.as_mask);
sigaction(<SIGNAL_NUMBER>, &siga, NULL);
```

The alternate stack is not inherited.

## Parameters

<i>thread</i>	Points to where the thread ID will be stored.
<i>attr</i>	Specifies the thread attributes object to use in creating the thread. If the value is <b>NULL</b> , the default attributes values will be used.
<i>start_routine</i>	Points to the routine to be executed by the thread.
<i>arg</i>	Points to the single argument to be passed to the <i>start_routine</i> routine.

## Return Values

If successful, the **pthread\_create** function returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_create** function will fail if:

<b>EAGAIN</b>	If WLM is running, the limit on the number of threads in the class may have been met.
<b>EINVAL</b>	The value specified by <b>attr</b> is invalid.
<b>EPERM</b>	The caller does not have appropriate permission to set the required scheduling parameters or scheduling policy.

The **pthread\_create** function will not return an error code of **EINTR**.

## Related Information

The **core** file format.

The **dbx** and **ulimit** commands.

The **pthread\_attr\_init** (“pthread\_attr\_init Subroutine” on page 1217) subroutine, **pthread\_attr\_destroy** (“pthread\_attr\_destroy Subroutine” on page 1209) subroutine, **pthread\_exit** (“pthread\_exit Subroutine” on page 1248) subroutine, **pthread\_cancel** (“pthread\_cancel Subroutine” on page 1231) subroutine, **pthread\_kill** (“pthread\_kill Subroutine” on page 1265) subroutine, **pthread\_self** (“pthread\_self Subroutine” on page 1295) subroutine, **pthread\_once** (“pthread\_once Subroutine” on page 1283) subroutine, **pthread\_join** (“pthread\_join or pthread\_detach Subroutine” on page 1262) subroutine, **fork** (“fork, f\_fork, or vfork Subroutine” on page 291) subroutine, and the **pthread.h** file.

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_create\_withcred\_np Subroutine

### Purpose

Creates a new thread with a new set of credentials, initializes its attributes, and makes it runnable.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
#include <sys/cred.h>
```

```
int pthread_create_withcred_np(pthread_t *thread, const pthread_attr_t *attr,
void *(*start_routine)(void),
void *arg, struct __pthrdscreds *credp)
```

### Description

The **pthread\_create\_withcred\_np** subroutine is exactly equivalent to the **pthread\_create** routine except that it allows the new thread to be created and start running with the credentials specified by the *credp* parameter. Only a process that has the credentials capability or is running with an effective user ID as the root user is allowed to modify its credentials using this routine.

The following credentials can be modified:

- Effective, real and saved user IDs
- Effective, real and saved group IDs
- Supplementary group IDs

**Note:** The administrator can set the lowest user ID value to which a process with credentials capability is allowed to switch its user IDs. A value of 0 can be specified for any of the preceding credentials to indicate that the thread should inherit that specific credential from its caller. The administrator can also set the lowest group ID to which a process with credentials capability is allowed to switch its group IDs.

The *\_\_pc\_flags* flag field in the *credp* parameter provides options to inherit credentials from the parent thread.

The newly created thread runs with per-thread credentials, and system calls like **getuid** or **getgid** returns the thread's credentials. Similarly, when a file is opened or a message is received, the thread's credentials will be used to determine whether the thread has privilege to execute the operation.

## Parameters

<i>thread</i>	Points to where the thread ID will be stored.
<i>attr</i>	Specifies the thread attributes object to use in creating the thread. If the value is NULL, the default attributes values will be used.
<i>start_routine</i>	Points to the routine to be executed by the thread.
<i>arg</i>	Points to the single argument to be passed to the <b>start_routine</b> routine.
<i>credp</i>	Points to a structure of type <b>__pthrdscreds</b> , which contains the credentials structure and the inheritance flags. If set to NULL, the <b>pthread_create_withcred_np</b> subroutine is the same as the <b>pthread_create</b> routine. The <b>__pc_cred</b> field indicates the credentials to be assigned to the new pthread. The <b>__pc_flags</b> field indicates which credentials, if any, are to be inherited from the parent thread. This field is constructed by logically OR'ing one or more of the following values: <b>PTHRDSCREDS_INHERIT_UIDS</b> Inherit user IDs from parent thread. <b>PTHRDSCREDS_INHERIT_GIDS</b> Inherit group IDs from parent thread. <b>PTHRDSCREDS_INHERIT_GSETS</b> Inherit the group sets from parent thread. <b>PTHRDSCREDS_INHERIT_PRIVILEGES</b> Inherit privileges from the parent thread.

## Security

Only a process that has the credentials capability or is running with an effective user ID (such as the root user) is allowed to modify its credentials using this routine.

## Return Values

If successful, the **pthread\_create\_withcred\_np** subroutine returns 0. Otherwise, an error number is returned to indicate the error.

## Error Codes

<b>EAGAIN</b>	If WLM is running, the limit on the number of threads in the class might have been met.
<b>EFAULT</b>	The <i>credp</i> parameter points to a location outside of the allocated address space of the process.
<b>EINVAL</b>	The credentials specified in the <i>credp</i> parameter are not valid.
<b>EPERM</b>	The caller does not have appropriate permission to set the credentials.

The **pthread\_create\_withcred\_np** subroutine will not return an error code of **EINTR**.

## Related Information

"pthread\_create Subroutine" on page 1243

---

## pthread\_delay\_np Subroutine

### Purpose

Causes a thread to wait for a specified period.

## Library

Threads Library (**libpthreads.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_delay_np ( interval )  
struct timespec *interval;
```

## Description

The **pthread\_delay\_np** subroutine causes the calling thread to delay execution for a specified period of elapsed wall clock time. The period of time the thread waits is at least as long as the number of seconds and nanoseconds specified in the *interval* parameter.

### Notes:

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.
2. The **pthread\_delay\_np** subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

## Parameters

*interval*            Points to the time structure specifying the wait period.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The **pthread\_delay\_np** subroutine is unsuccessful if the following is true:

**EINVAL**        The *interval* parameter is not valid.

## Related Information

The **sleep**, **nsleep**, or **usleep** subroutine.

---

## pthread\_equal Subroutine

### Purpose

Compares two thread IDs.

### Library

Threads Library (**libpthreads.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_equal (thread1, thread2)
pthread_t thread1;
pthread_t thread2;
```

## Description

The **pthread\_equal** subroutine compares the thread IDs *thread1* and *thread2*. Since the thread IDs are opaque objects, it should not be assumed that they can be compared using the equality operator (==).

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

## Parameters

*thread1*        Specifies the first ID to be compared.  
*thread2*        Specifies the second ID to be compared.

## Return Values

The **pthread\_equal** function returns a nonzero value if *thread1* and *thread2* are equal; otherwise, zero is returned.

If either *thread1* or *thread2* are not valid thread IDs, the behavior is undefined.

## Related Information

The **pthread\_self** (“pthread\_self Subroutine” on page 1295) subroutine, the **pthread\_create** (“pthread\_create Subroutine” on page 1243) subroutine, the **pthread.h** file.

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_exit Subroutine

### Purpose

Terminates the calling thread.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
void pthread_exit (status)
void *status;
```

### Description

The **pthread\_exit** subroutine terminates the calling thread safely, and stores a termination status for any thread that may join the calling thread. The termination status is always a void pointer; it can reference any kind of data. It is not recommended to cast this pointer into a scalar data type (**int** for example), because the casts may not be portable. This subroutine never returns.



Unlike the **exit** subroutine, the **pthread\_exit** subroutine does not close files. Thus any file opened and used only by the calling thread must be closed before calling this subroutine. It is also important to note that the **pthread\_exit** subroutine frees any thread-specific data, including the thread's stack. Any data allocated on the stack becomes invalid, since the stack is freed and the corresponding memory may be reused by another thread. Therefore, thread synchronization objects (mutexes and condition variables) allocated on a thread's stack must be destroyed before the thread calls the **pthread\_exit** subroutine.

Returning from the initial routine of a thread implicitly calls the **pthread\_exit** subroutine, using the return value as parameter.

If the thread is not detached, its resources, including the thread ID, the termination status, the thread-specific data, and its storage, are all maintained until the thread is detached or the process terminates.

If another thread joins the calling thread, that thread wakes up immediately, and the calling thread is automatically detached.

If the thread is detached, the cleanup routines are popped from their stack and executed. Then the destructor routines from the thread-specific data are executed. Finally, the storage of the thread is reclaimed and its ID is freed for reuse.

Terminating the initial thread by calling this subroutine does not terminate the process, it just terminates the initial thread. However, if all the threads in the process are terminated, the process is terminated by implicitly calling the **exit** subroutine with a return code of 0 if the last thread is detached, or 1 otherwise.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

## Parameters

*status* Points to an optional termination status, used by joining threads. If no termination status is desired, its value should be **NULL**.

## Return Values

The **pthread\_exit** function cannot return to its caller.

## Errors

No errors are defined.

The **pthread\_exit** function will not return an error code of **EINTR**.

## Related Information

The **pthread\_cleanup\_push** (“**pthread\_cleanup\_pop** or **pthread\_cleanup\_push** Subroutine” on page 1232) subroutine, **pthread\_cleanup\_pop** (“**pthread\_cleanup\_pop** or **pthread\_cleanup\_push** Subroutine” on page 1232) subroutine, **pthread\_key\_create** (“**pthread\_key\_create** Subroutine” on page 1263) subroutine, **pthread\_create** (“**pthread\_create** Subroutine” on page 1243) subroutine, **pthread\_join** (“**pthread\_join** or **pthread\_detach** Subroutine” on page 1262) subroutine, **pthread\_cancel** (“**pthread\_cancel** Subroutine” on page 1231) subroutine, **exit** (“**exit**, **atexit**, **unatexit**, **\_exit**, or **\_Exit** Subroutine” on page 243) subroutine, the **pthread.h** file.

Terminating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_get\_expiration\_np Subroutine

### Purpose

Obtains a value representing a desired expiration time.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_get_expiration_np ( delta, abstime )  
struct timespec *delta;  
struct timespec *abstime;
```

### Description

The **pthread\_get\_expiration\_np** subroutine adds the interval *delta* to the current absolute system time and returns a new absolute time. This new absolute time can be used as the expiration time in a call to the **pthread\_cond\_timedwait** subroutine.

#### Notes:

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.
2. The **pthread\_get\_expiration\_np** subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

### Parameters

*delta*            Points to the time structure specifying the interval.  
*abstime*        Points to where the new absolute time will be stored.

### Return Values

Upon successful completion, the new absolute time is returned via the *abstime* parameter, and 0 is returned. Otherwise, an error code is returned.

### Error Codes

The **pthread\_get\_expiration\_np** subroutine is unsuccessful if the following is true:

**EINVAL**        The *delta* or *abstime* parameters are not valid.

### Related Information

The **pthread\_cond\_timedwait** (“**pthread\_cond\_wait** or **pthread\_cond\_timedwait** Subroutine” on page 1236) subroutine.

---

## pthread\_getconcurrency or pthread\_setconcurrency Subroutine

### Purpose

Gets or sets level of concurrency.

### Library

Threads Library (**libthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_getconcurrency (void);
```

```
int pthread_setconcurrency (new_level)  
int new_level;
```

### Description

The **pthread\_setconcurrency** subroutine allows an application to inform the threads implementation of its desired concurrency level, *new\_level*. The actual level of concurrency provided by the implementation as a result of this function call is unspecified.

If *new\_level* is zero, it causes the implementation to maintain the concurrency level at its discretion as if **pthread\_setconcurrency** was never called.

The **pthread\_getconcurrency** subroutine returns the value set by a previous call to the **pthread\_setconcurrency** subroutine. If the **pthread\_setconcurrency** subroutine was not previously called, this function returns zero to indicate that the implementation is maintaining the concurrency level.

When an application calls **pthread\_setconcurrency**, it is informing the implementation of its desired concurrency level. The implementation uses this as a hint, not a requirement.

Use of these subroutines changes the state of the underlying concurrency upon which the application depends. Library developers are advised to not use the **pthread\_getconcurrency** and **pthread\_setconcurrency** subroutines since their use may conflict with an applications use of these functions.

### Parameters

*new\_level*                                  Specifies the value of the concurrency level.

### Return Value

If successful, the **pthread\_setconcurrency** subroutine returns zero. Otherwise, an error number is returned to indicate the error.

The **pthread\_getconcurrency** subroutine always returns the concurrency level set by a previous call to **pthread\_setconcurrency**. If the **pthread\_setconcurrency** subroutine has never been called, **pthread\_getconcurrency** returns zero.

### Error Codes

The **pthread\_setconcurrency** subroutine will fail if:

**EINVAL**        The value specified by *new\_level* is negative.

**EAGAIN** The value specific by *new\_level* would cause a system resource to be exceeded.

## Related Information

The `pthread.h` file.

---

## pthread\_getcpuclockid Subroutine

### Purpose

Accesses a thread CPU-time clock.

### Syntax

```
#include <pthread.h>
#include <time.h>
```

```
int pthread_getcpuclockid(pthread_t thread_id, clockid_t *clock_id);
```

### Description

The `pthread_getcpuclockid` subroutine returns in the *clock\_id* parameter the clock ID of the CPU-time clock of the thread specified by *thread\_id*, if the thread specified by *thread\_id* exists.

### Parameters

<i>thread_id</i>	Specifies the ID of the pthread whose clock ID is requested.
<i>clock_id</i>	Points to the <code>clockid_t</code> structure used to return the thread CPU-time clock ID of <i>thread_id</i> .

### Return Values

Upon successful completion, the `pthread_getcpuclockid` subroutine returns 0; otherwise, an error number is returned to indicate the error.

### Error Codes

<b>ENOTSUP</b>	The subroutine is not supported with checkpoint-restart'ed processes.
<b>ESRCH</b>	The value specified by <i>thread_id</i> does not refer to an existing thread.

## Related Information

“clock\_getcpuclockid Subroutine” on page 169, “clock\_getres, clock\_gettime, and clock\_settime Subroutine” on page 170, timer\_create Subroutine, timer\_gettime Subroutine

---

## pthread\_getrusage\_np Subroutine

### Purpose

Enable or disable pthread library resource collection, and retrieve resource information for any pthread in the current process.

### Library

Threads Library (`libpthreads.a`)

## Syntax

```
#include <pthread.h>
```

```
int pthread_getrusage_np (Ptid, RUsage, Mode)  
pthread_t Ptid;  
struct rusage *RUsage;  
int Mode;
```

## Description

The **pthread\_getrusage\_np** subroutine enables and disables resource collection in the pthread library and collects resource information for any pthread in the current process. When compiled in 64-bit mode, resource usage (rusage) counters are 64-bits for the calling thread. When compiled in 32-bit mode, rusage counters are 32-bits for the calling pthread.

This functionality is enabled by default. The previous **AIXTHREAD\_ENRUSG** used with **pthread\_getrusage\_np** is no longer supported.

## Parameters

*Ptid*                Specifies the target thread. Must be within the current process.

## *RUsage*

Points to a buffer described in the `/usr/include/sys/resource.h` file. The fields are defined as follows:

### **ru\_utime**

The total amount of time running in user mode.

### **ru\_stime**

The total amount of time spent in the system executing on behalf of the processes.

### **ru\_maxrss**

The maximum size, in kilobytes, of the used resident set size.

### **ru\_ixrss**

An integral value indicating the amount of memory used by the text segment that was also shared among other processes. This value is expressed in *units of kilobytes X seconds-of-execution* and is calculated by adding the number of shared memory pages in use each time the internal system clock ticks, and then averaging over one-second intervals.

### **ru\_idrss**

An integral value of the amount of unshared memory in the data segment of a process, which is expressed in *units of kilobytes X seconds-of-execution*.

### **ru\_minflt**

The number of page faults serviced without any I/O activity. In this case, I/O activity is avoided by reclaiming a page frame from the list of pages awaiting reallocation.

### **ru\_majflt**

The number of page faults serviced that required I/O activity.

### **ru\_nswap**

The number of times that a process was swapped out of main memory.

### **ru\_inblock**

The number of times that the file system performed input.

### **ru\_oublock**

The number of times that the file system performed output.

**Note:** The numbers that the `ru_inblock` and `ru_oublock` fields display account for real I/O only; data supplied by the caching mechanism is charged only to the first process that reads or writes the data.

### **ru\_msgsnd**

The number of IPC messages sent.

### **ru\_msgrcv**

The number of IPC messages received.

### **ru\_nsignals**

The number of signals delivered.

### **ru\_nvcsw**

The number of times a context switch resulted because a process voluntarily gave up the processor before its time slice was completed. This usually occurs while the process waits for a resource to become available.

### **ru\_nivcsw**

The number of times a context switch resulted because a higher priority process ran or because the current process exceeded its time slice.

*Mode* Indicates which task the subroutine should perform. Acceptable values are as follows:

**PTHRDSINFO\_RUSAGE\_START**

Returns the current resource utilization, which will be the start measurement.

**PTHRDSINFO\_RUSAGE\_STOP**

Returns total current resource utilization since the last time a **PTHRDSINFO\_RUSAGE\_START** was performed. If the task **PTHRDSINFO\_RUSAGE\_START** was not performed, then the resource information returned is the accumulated value since the start of the pthread.

**PTHRDSINFO\_RUSAGE\_COLLECT**

Collects resource information for the target thread. If the task **PTHRDSINFO\_RUSAGE\_START** was not performed, then the resource information returned is the accumulated value since the start of the pthread.

## Return Values

Upon successful completion, the **pthread\_getrusage\_np** subroutine returns a value of 0. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_getrusage\_np** subroutine fails if:

**EINVAL** The address specified for *RUsage* is NULL, not valid, or a null value for *Ptid* was given.

**ESRCH** Either no thread could be found corresponding to the ID thread of the *Ptid* thread or the thread corresponding to the *Ptid* thread ID was not in the current process.

## Related Information

The **pthread.h** subroutine.

---

## pthread\_getschedparam Subroutine

### Purpose

Returns the current schedpolicy and schedparam attributes of a thread.

### Library

Threads Library (**libpthread.a**)

### Syntax

```
#include <pthread.h>
#include <sys/sched.h>
```

```
int pthread_getschedparam ( thread, schedpolicy, schedparam)
pthread_t thread;
int *schedpolicy;
struct sched_param *schedparam;
```

### Description

The **pthread\_getschedparam** subroutine returns the current schedpolicy and schedparam attributes of the thread *thread*. The schedpolicy attribute specifies the scheduling policy of a thread. It may have one of the following values:

**SCHED\_FIFO** Denotes first-in first-out scheduling.

<b>SCHED_RR</b>	Denotes round-robin scheduling.
<b>SCHED_OTHER</b>	Denotes the default operating system scheduling policy. It is the default value.

The schedparam attribute specifies the scheduling parameters of a thread created with this attributes object. The sched\_priority field of the **sched\_param** structure contains the priority of the thread. It is an integer value.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the cc\_r compiler used. In this case, the flag is automatically set.

The implementation of this subroutine is dependent on the priority scheduling POSIX option. The priority scheduling POSIX option is implemented in the operating system.

## Parameters

<i>thread</i>	Specifies the target thread.
<i>schedpolicy</i>	Points to where the schedpolicy attribute value will be stored.
<i>schedparam</i>	Points to where the schedparam attribute value will be stored.

## Return Values

Upon successful completion, the current value of the schedpolicy and schedparam attributes are returned via the *schedpolicy* and *schedparam* parameters, and 0 is returned. Otherwise, an error code is returned.

## Error Codes

The **pthread\_getschedparam** subroutine is unsuccessful if the following is true:

**ESRCH** The thread *thread* does not exist.

## Related Information

The **pthread\_attr\_getschedparam** (“pthread\_attr\_getschedparam Subroutine” on page 1213) subroutine.

Threads Scheduling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_getspecific or pthread\_setspecific Subroutine

### Purpose

Returns and sets the thread-specific data associated with the specified key.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
void *pthread_getspecific (key)
pthread_key_t key;
```



```
int pthread_setspecific (key, value)
pthread_key_t key;
const void *value;
```

## Description

The **pthread\_setspecific** function associates a thread-specific *value* with a *key* obtained via a previous call to **pthread\_key\_create**. Different threads may bind different values to the same key. These values are typically pointers to blocks of dynamically allocated memory that have been reserved for use by the calling thread.

The **pthread\_getspecific** function returns the value currently bound to the specified *key* on behalf of the calling thread.

The effect of calling **pthread\_setspecific** or **pthread\_getspecific** with a *key* value not obtained from **pthread\_key\_create** or after key has been deleted with **pthread\_key\_delete** is undefined.

Both **pthread\_setspecific** and **pthread\_getspecific** may be called from a thread-specific data destructor function. However, calling **pthread\_setspecific** from a destructor may result in lost storage or infinite loops.

## Parameters

*key* Specifies the key to which the value is bound.  
*value* Specifies the new thread-specific value.

## Return Values

The function **pthread\_getspecific** returns the thread-specific data value associated with the given key. If no thread-specific data value is associated with key, then the value NULL is returned. If successful, the **pthread\_setspecific** function returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_setspecific** function will fail if:

**ENOMEM** Insufficient memory exists to associate the value with the key.

The **pthread\_setspecific** function may fail if:

**EINVAL** The key value is invalid.

No errors are returned from **pthread\_getspecific**.

These functions will not return an error code of EINTR.

## Related Information

The **pthread\_key\_create** (“pthread\_key\_create Subroutine” on page 1263) subroutine, the **pthread.h** file.

Thread-Specific Data in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_getthrds\_np Subroutine

### Purpose

Retrieves register and stack information for threads.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_getthrds_np (thread, mode, buf, bufsize, regbuf, regbufsize)
pthread_t *ptid;
int mode;
struct __pthrdsinfo *buf;
int bufsize;
void *regbuf;
int *regbufsize;
```

### Description

The **pthread\_getthrds\_np** subroutine retrieves information on the state of the thread *thread* and its underlying kernel thread, including register and stack information.

### Parameters

*thread*            The pointer to the thread. On input it identifies the target thread of the operation, or 0 to operate on the first entry in the list of threads. On output it identifies the next entry in the list of threads, or 0 if the end of the list has been reached. **pthread\_getthrds\_np** can be used to traverse the whole list of threads by starting with *thread* pointing to 0 and calling **pthread\_getthrds\_np** repeatedly until it returns with *thread* pointing to 0.

*mode*

Specifies the type of query. These values can be bitwise or'ed together to specify more than one type of query.

**PTHRDSINFO\_QUERY\_GPRS**

get general purpose registers

**PTHRDSINFO\_QUERY\_SPRS**

get special purpose registers

**PTHRDSINFO\_QUERY\_FPRS**

get floating point registers

**PTHRDSINFO\_QUERY\_REGS**

get all of the above registers

**PTHRDSINFO\_QUERY\_TID**

get the kernel thread id

**PTHRDSINFO\_QUERY\_TLS**

get the thread-local storage information.

This value can be or'ed with any value of the mode parameter. The thread-local storage information is returned to the caller in a caller-provided buffer, **regbuf**. If the buffer is too small for the data, the buffer is filled up to the end of the buffer and *ERANGE* is returned. The caller also provides the size of the buffer, **regbufsize**, which on return is changed to the size of the thread local storage information even if it does not fit into a buffer.

The thread-local storage information is returned in form of an array of tuples: memory address and TLS region (unique number assigned by the loader). The TLS region is also included in the loader info structure returned by **loadquery**. If you need any additional information such as TLS size, you can find it in that structure.

```
#typedef struct __pthrdstlsinfo{
    void *pti_vaddr;
    int   pti_region;
} PTHRDS_TLS_INFO;
```

**PTHRDSINFO\_QUERY\_EXTCTX**

get the extended machine context

**PTHRDSINFO\_QUERY\_ALL**

get everything (except for the extended context, which must be explicitly requested)

*buf*

Specifies the address of the struct **\_\_pthrdsinfo** structure that will be filled in by **pthread\_getthrds\_np**. On return, this structure holds the following data (depending on the type of query requested):

**\_\_pi\_ptid**

The thread's thread identifier

**\_\_pi\_tid**

The thread's kernel thread id, or 0 if the thread does not have a kernel thread

**\_\_pi\_state**

The state of the thread, equal to one of the following:

**PTHRDSINFO\_STATE\_RUN**

The thread is running

**PTHRDSINFO\_STATE\_READY**

The thread is ready to run

**PTHRDSINFO\_STATE\_IDLE**

The thread is being initialized

**PTHRDSINFO\_STATE\_SLEEP**

The thread is sleeping

**PTHRDSINFO\_STATE\_TERM**

The thread is terminated

**PTHRDSINFO\_STATE\_NOTSUP**

Error condition

**\_\_pi\_suspended**

1 if the thread is suspended, 0 if it is not

**\_\_pi\_returned**

The return status of the thread

**\_\_pi\_ustk**

The thread's user stack pointer

**\_\_pi\_context**

The thread's context (register information)

If the **PTHRDSINFO\_QUERY\_EXTCTX** mode is requested, then the *buf* specifies the address of a **\_\_pthrdsinfox** structure, which, in addition to all of the preceding information, also contains the following:

**\_\_pi\_ec**

The thread's extended context (extended register state)

*bufsize*

The size of the **\_\_pthrdsinfo** or **\_\_pthrdsinfox** structure in bytes.

*regbuf*

The location of the buffer to hold the register save data and to pass the TLS information from the kernel if the thread is in a system call.

*regbufsize*

The pointer to the size of the *regbuf* buffer. On input, it identifies the maximum size of the buffer in bytes. On output, it identifies the number of bytes of register save data and pass the TLS information. If the thread is not in a system call, there is no register save data returned from the kernel, and *regbufsize* is 0. If the size of the register save data is larger than the input value of *regbufsize*, the number of bytes specified by the input value of *regbufsize* is copied to *regbuf*, **pthread\_getthrds\_np()** returns ERANGE, and the output value of *regbufsize* specifies the number of bytes required to hold all of the register save data.

## Return Values

If successful, the **pthread\_getthrds\_np** function returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_getthrds_np` function will fail if:

- EINVAL** Either *thread* or *buf* is NULL, or *bufsize* is not equal to the size of the `__pthrdsinfo` structure in the library.
- ESRCH** No thread could be found corresponding to that specified by the thread ID *thread*.
- ERANGE** *regbuf* was not large enough to handle all of the register save data.
- ENOMEM** Insufficient memory exists to perform this operation.

## Related Information

The `pthread.h` file.

---

## pthread\_getunique\_np Subroutine

### Purpose

Returns the sequence number of a thread.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_getunique_np ( thread, sequence)
pthread_t *thread;
int *sequence;
```

### Description

The `pthread_getunique_np` subroutine returns the sequence number of the thread *thread*. The sequence number is a number, unique to each thread, associated with the thread at creation time.

#### Notes:

1. The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.
2. The `pthread_getunique_np` subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

### Parameters

- thread* Specifies the thread.
- sequence* Points to where the sequence number will be stored.

### Return Values

Upon successful completion, the sequence number is returned via the *sequence* parameter, and 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_getunique_np` subroutine is unsuccessful if the following is true:

<b>EINVAL</b>	The <i>thread</i> or <i>sequence</i> parameters are not valid.
<b>ESRCH</b>	The thread <i>thread</i> does not exist.

## Related Information

The `pthread_self` (“`pthread_self` Subroutine” on page 1295) subroutine.

---

## pthread\_join or pthread\_detach Subroutine

### Purpose

Blocks or detaches the calling thread until the specified thread terminates.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_join (thread, status)
pthread_t thread;
void **status;
```

```
int pthread_detach (thread)
pthread_t thread;
```

### Description

The `pthread_join` subroutine blocks the calling thread until the thread *thread* terminates. The target thread's termination status is returned in the *status* parameter.

If the target thread is already terminated, but not yet detached, the subroutine returns immediately. It is impossible to join a detached thread, even if it is not yet terminated. The target thread is automatically detached after all joined threads have been woken up.

This subroutine does not itself cause a thread to be terminated. It acts like the `pthread_cond_wait` subroutine to wait for a special condition.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

The `pthread_detach` subroutine is used to indicate to the implementation that storage for the thread whose thread ID is in the location *thread* can be reclaimed when that thread terminates. This storage shall be reclaimed on process exit, regardless of whether the thread has been detached or not, and may include storage for *thread* return value. If *thread* has not yet terminated, `pthread_detach` shall not cause it to terminate. Multiple `pthread_detach` calls on the same target thread causes an error.

### Parameters

<i>thread</i>	Specifies the target thread.
<i>status</i>	Points to where the termination status of the target thread will be stored. If the value is <b>NULL</b> , the termination status is not returned.

## Return Values

If successful, the `pthread_join` function returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_join` and `pthread_detach` functions will fail if:

**EINVAL** The implementation has detected that the value specified by thread does not refer to a joinable thread.  
**ESRCH** No thread could be found corresponding to that specified by the given thread ID.

The `pthread_join` function will fail if:

**EDEADLK** The value of thread specifies the calling thread.

The `pthread_join` function will not return an error code of **EINTR**.

## Related Information

The `pthread_exit` (“`pthread_exit` Subroutine” on page 1248) subroutine, `pthread_create` (“`pthread_create` Subroutine” on page 1243) subroutine, `wait` subroutine, `pthread_cond_wait` or `pthread_cond_timedwait` (“`pthread_cond_wait` or `pthread_cond_timedwait` Subroutine” on page 1236) subroutines, the `pthread.h` file.

Joining Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_key\_create Subroutine

### Purpose

Creates a thread-specific data key.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_key_create ( key, destructor )  
pthread_key_t * key;  
void (* destructor) (void *);
```

### Description

The `pthread_key_create` subroutine creates a thread-specific data key. The key is shared among all threads within the process, but each thread has specific data associated with the key. The thread-specific data is a void pointer, initially set to **NULL**.

The application is responsible for ensuring that this subroutine is called only once for each requested key. This can be done, for example, by calling the subroutine before creating other threads, or by using the one-time initialization facility.

Typically, thread-specific data are pointers to dynamically allocated storage. When freeing the storage, the value should be set to **NULL**. It is not recommended to cast this pointer into scalar data type (**int** for example), because the casts may not be portable, and because the value of **NULL** is implementation dependent.

An optional destructor routine can be specified. It will be called for each thread when it is terminated and detached, after the call to the cleanup routines, if the specific value is not **NULL**. Typically, the destructor routine will release the storage thread-specific data. It will receive the thread-specific data as a parameter.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

## Parameters

<i>key</i>	Points to where the key will be stored.
<i>destructor</i>	Points to an optional destructor routine, used to cleanup data on thread termination. If no cleanup is desired, this pointer should be <b>NULL</b> .

## Return Values

If successful, the **pthread\_key\_create** function stores the newly created key value at *\*key* and returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_key\_create** function will fail if:

<b>EAGAIN</b>	The system lacked the necessary resources to create another thread-specific data key, or the system-imposed limit on the total number of keys per process <b>PTHREAD_KEYS_MAX</b> has been exceeded.
<b>ENOMEM</b>	Insufficient memory exists to create the key.

The **pthread\_key\_create** function will not return an error code of **EINTR**.

## Related Information

The **pthread\_exit** (“pthread\_exit Subroutine” on page 1248) subroutine, **pthread\_key\_delete** (“pthread\_key\_delete Subroutine”) subroutine, **pthread\_getspecific** (“pthread\_getspecific or pthread\_setspecific Subroutine” on page 1256) subroutine, **pthread\_once** (“pthread\_once Subroutine” on page 1283) subroutine, **pthread.h** file.

Thread-Specific Data in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_key\_delete Subroutine

### Purpose

Deletes a thread-specific data key.

### Library

Threads Library (**libpthreads.a**)



## Syntax

```
#include <pthread.h>
```

```
int pthread_key_delete (key)
pthread_key_t key;
```

## Description

The **pthread\_key\_delete** subroutine deletes the thread-specific data key *key*, previously created with the **pthread\_key\_create** subroutine. The application must ensure that no thread-specific data is associated with the key. No destructor routine is called.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

## Parameters

*key* Specifies the key to delete.

## Return Values

If successful, the **pthread\_key\_delete** function returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_key\_delete** function will fail if:

**EINVAL** The key value is invalid.

The **pthread\_key\_delete** function will not return an error code of **EINTR**.

## Related Information

The **pthread\_key\_create** (“pthread\_key\_create Subroutine” on page 1263) subroutine, **pthread.h** file.

Thread-Specific Data in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_kill Subroutine

### Purpose

Sends a signal to the specified thread.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <signal.h>
```

```
int pthread_kill (thread, signal)
pthread_t thread;
int signal;
```

## Description

The `pthread_kill` subroutine sends the signal *signal* to the thread *thread*. It acts with threads like the `kill` subroutine with single-threaded processes.

If the receiving thread has blocked delivery of the signal, the signal remains pending on the thread until the thread unblocks delivery of the signal or the action associated with the signal is set to ignore the signal.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

## Parameters

*thread* Specifies the target thread for the signal.  
*signal* Specifies the signal to be delivered. If the signal value is 0, error checking is performed, but no signal is delivered.

## Return Values

Upon successful completion, the function returns a value of zero. Otherwise the function returns an error number. If the `pthread_kill` function fails, no signal is sent.

## Error Codes

The `pthread_kill` function will fail if:

**ESRCH** No thread could be found corresponding to that specified by the given thread ID.  
**EINVAL** The value of the *signal* parameter is an invalid or unsupported signal number.

The `pthread_kill` function will not return an error code of **EINTR**.

## Related Information

The `kill` (“kill or killpg Subroutine” on page 580) subroutine, `pthread_cancel` (“pthread\_cancel Subroutine” on page 1231) subroutine, `pthread_create` (“pthread\_create Subroutine” on page 1243) subroutine, `sigaction` subroutine, `pthread_self` (“pthread\_self Subroutine” on page 1295) subroutine, `raise` subroutine, `pthread.h` file.

Signal Management in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_lock\_global\_np Subroutine

### Purpose

Locks the global mutex.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
void pthread_lock_global_np ()
```

## Description

The **pthread\_lock\_global\_np** subroutine locks the global mutex. If the global mutex is currently held by another thread, the calling thread waits until the global mutex is unlocked. The subroutine returns with the global mutex locked by the calling thread.

Use the global mutex when calling a library package that is not designed to run in a multithreaded environment. (Unless the documentation for a library function specifically states that it is compatible with multithreading, assume that it is not compatible; in other words, assume it is nonreentrant.)

The global mutex is one lock. Any code that calls any function that is not known to be reentrant uses the same lock. This prevents dependencies among threads calling library functions and those functions calling other functions, and so on.

The global mutex is a recursive mutex. A thread that has locked the global mutex can relock it without deadlocking. The thread must then call the **pthread\_unlock\_global\_np** subroutine as many times as it called this routine to allow another thread to lock the global mutex.

### Notes:

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.
2. The **pthread\_lock\_global\_np** subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

## Related Information

The **pthread\_mutex\_lock** (“pthread\_mutex\_lock, pthread\_mutex\_trylock, or pthread\_mutex\_unlock Subroutine” on page 1270) subroutine, **pthread\_unlock\_global\_np** (“pthread\_unlock\_global\_np Subroutine” on page 1305) subroutine.

Using Mutexes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_mutex\_init or pthread\_mutex\_destroy Subroutine

### Purpose

Initializes or destroys a mutex.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_mutex_init (mutex, attr)
pthread_mutex_t *mutex;
const pthread_mutexattr_t *attr;
```

```
int pthread_mutex_destroy (mutex)
pthread_mutex_t *mutex;
```

## Description

The **pthread\_mutex\_init** function initializes the mutex referenced by *mutex* with attributes specified by *attr*. If *attr* is NULL, the default mutex attributes are used; the effect is the same as passing the address of a default mutex attributes object. Upon successful initialization, the state of the mutex becomes initialized and unlocked.

Attempting to initialize an already initialized mutex results in undefined behavior.

The **pthread\_mutex\_destroy** function destroys the mutex object referenced by *mutex*; the mutex object becomes, in effect, uninitialized. An implementation may cause **pthread\_mutex\_destroy** to set the object referenced by *mutex* to an invalid value. A destroyed mutex object can be re-initialized using **pthread\_mutex\_init**; the results of otherwise referencing the object after it has been destroyed are undefined.

It is safe to destroy an initialized mutex that is unlocked. Attempting to destroy a locked mutex results in undefined behavior.

In cases where default mutex attributes are appropriate, the macro **PTHREAD\_MUTEX\_INITIALIZER** can be used to initialize mutexes that are statically allocated. The effect is equivalent to dynamic initialization by a call to **pthread\_mutex\_init** with parameter *attr* specified as NULL, except that no error checks are performed.

## Parameters

*mutex*                Specifies the mutex to initialize or delete.  
*attr*                 Specifies the mutex attributes object.

## Return Values

If successful, the **pthread\_mutex\_init** and **pthread\_mutex\_destroy** functions return zero. Otherwise, an error number is returned to indicate the error. The EBUSY and EINVAL error checks act as if they were performed immediately at the beginning of processing for the function and cause an error return prior to modifying the state of the mutex specified by *mutex*.

## Error Codes

The **pthread\_mutex\_init** function will fail if:

**ENOMEM**    Insufficient memory exists to initialize the mutex.  
**EINVAL**     The value specified by *attr* is invalid.  
**EPERM**     The caller does not have the privilege to perform the operation in a strictly standards conforming environment where environment variable **XPG\_SUS\_ENV=ON**.

The **pthread\_mutex\_destroy** function may fail if:

**EBUSY**     The implementation has detected an attempt to destroy the object referenced by *mutex* while it is locked or referenced (for example, while being used in a **pthread\_cond\_wait** or **pthread\_cond\_timedwait** by another thread).  
**EINVAL**     The value specified by *mutex* is invalid.

These functions will not return an error code of EINTR.

## Related Information

The **pthread\_mutex\_lock**, **pthread\_mutex\_trylock** (“pthread\_mutex\_lock, pthread\_mutex\_trylock, or pthread\_mutex\_unlock Subroutine” on page 1270) subroutine, **pthread\_mutex\_unlock** (“pthread\_mutex\_lock, pthread\_mutex\_trylock, or pthread\_mutex\_unlock Subroutine” on page 1270) subroutine, **pthread\_mutexattr\_setpshared** (“pthread\_mutexattr\_getpshared or pthread\_mutexattr\_setpshared Subroutine” on page 1279) subroutine.

The **pthread.h** file.

---

## pthread\_mutex\_getprioceiling or pthread\_mutex\_setprioceiling Subroutine

### Purpose

Gets and sets the priority ceiling of a mutex.

### Syntax

```
#include <pthread.h>
```

```
int pthread_mutex_getprioceiling(const pthread_mutex_t *restrict mutex,
    int *restrict prioceiling);
int pthread_mutex_setprioceiling(pthread_mutex_t *restrict mutex,
    int prioceiling, int *restrict old_ceiling);
```

### Description

The **pthread\_mutex\_getprioceiling** subroutine returns the current priority ceiling of the mutex.

The **pthread\_mutex\_setprioceiling** subroutine either locks the mutex if it is unlocked, or blocks until it can successfully lock the mutex, then it changes the mutex’s priority ceiling and releases the mutex. When the change is successful, the previous value of the priority ceiling shall be returned in *old\_ceiling*. The process of locking the mutex need not adhere to the priority protect protocol.

If the **pthread\_mutex\_setprioceiling** subroutine fails, the mutex priority ceiling is not changed.

### Return Values

If successful, the **pthread\_mutex\_getprioceiling** and **pthread\_mutex\_setprioceiling** subroutines return zero; otherwise, an error number is returned to indicate the error.

### Error Codes

The **pthread\_mutex\_getprioceiling** and **pthread\_mutex\_setprioceiling** subroutines can fail if:

<b>EINVAL</b>	The priority requested by the <i>prioceiling</i> parameter is out of range.
<b>EINVAL</b>	The value specified by the <i>mutex</i> parameter does not refer to a currently existing mutex.
<b>ENOSYS</b>	This function is not supported (draft 7).
<b>ENOTSUP</b>	This function is not supported together with checkpoint/restart.
<b>EPERM</b>	The caller does not have the privilege to perform the operation in a strictly standards conforming environment where environment variable <b>XPG_SUS_ENV=ON</b> .

## Related Information

The “pthread\_mutex\_init or pthread\_mutex\_destroy Subroutine” on page 1267, “pthread\_mutex\_lock, pthread\_mutex\_trylock, or pthread\_mutex\_unlock Subroutine” on page 1270, “pthread\_mutex\_timedlock Subroutine” on page 1272.

The `pthread.h` file.

---

## PTHREAD\_MUTEX\_INITIALIZER Macro

### Purpose

Initializes a static mutex with default attributes.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
static pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
```

### Description

The `PTHREAD_MUTEX_INITIALIZER` macro initializes the static mutex *mutex*, setting its attributes to default values. This macro should only be used for static mutexes, as no error checking is performed.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

### Related Information

The `pthread_mutex_init` (“`pthread_mutex_init` or `pthread_mutex_destroy` Subroutine” on page 1267) subroutine.

Using Mutexes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_mutex\_lock, pthread\_mutex\_trylock, or pthread\_mutex\_unlock Subroutine

### Purpose

Locks and unlocks a mutex.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>

int pthread_mutex_lock ( mutex )
pthread_mutex_t *mutex;

int pthread_mutex_trylock ( mutex )
pthread_mutex_t *mutex;

int pthread_mutex_unlock ( mutex )
pthread_mutex_t *mutex;
```

## Description

The mutex object referenced by the *mutex* parameter is locked by calling **pthread\_mutex\_lock**. If the mutex is already locked, the calling thread blocks until the mutex becomes available. This operation returns with the mutex object referenced by the *mutex* parameter in the locked state with the calling thread as its owner.

If the mutex type is `PTHREAD_MUTEX_NORMAL`, deadlock detection is not provided. Attempting to relock the mutex causes deadlock. If a thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, undefined behavior results.

If the mutex type is `PTHREAD_MUTEX_ERRORCHECK`, then error checking is provided. If a thread attempts to relock a mutex that it has already locked, an error will be returned. If a thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error will be returned.

If the mutex type is `PTHREAD_MUTEX_RECURSIVE`, then the mutex maintains the concept of a lock count. When a thread successfully acquires a mutex for the first time, the lock count is set to one. Each time the thread relocks this mutex, the lock count is incremented by one. Each time the thread unlocks the mutex, the lock count is decremented by one. When the lock count reaches zero, the mutex becomes available for other threads to acquire. If a thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error will be returned.

If the mutex type is `PTHREAD_MUTEX_DEFAULT`, attempting to recursively lock the mutex results in undefined behavior. Attempting to unlock the mutex if it was not locked by the calling thread results in undefined behavior. Attempting to unlock the mutex if it is not locked results in undefined behavior.

The function **pthread\_mutex\_trylock** is identical to **pthread\_mutex\_lock** except that if the mutex object referenced by the *mutex* parameter is currently locked (by any thread, including the current thread), the call returns immediately.

The **pthread\_mutex\_unlock** function releases the mutex object referenced by *mutex*. The manner in which a mutex is released is dependent upon the mutex's type attribute. If there are threads blocked on the mutex object referenced by the *mutex* parameter when **pthread\_mutex\_unlock** is called, resulting in the mutex becoming available, the scheduling policy is used to determine which thread will acquire the mutex. (In the case of `PTHREAD_MUTEX_RECURSIVE` mutexes, the mutex becomes available when the count reaches zero and the calling thread no longer has any locks on this mutex).

If a signal is delivered to a thread waiting for a mutex, upon return from the signal handler the thread resumes waiting for the mutex as if it was not interrupted.

## Parameter

*mutex* Specifies the mutex to lock.

## Return Values

If successful, the **pthread\_mutex\_lock** and **pthread\_mutex\_unlock** functions return zero. Otherwise, an error number is returned to indicate the error.

The function **pthread\_mutex\_trylock** returns zero if a lock on the mutex object referenced by the *mutex* parameter is acquired. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_mutex_trylock` function will fail if:

**EBUSY** The mutex could not be acquired because it was already locked.

The `pthread_mutex_lock`, `pthread_mutex_trylock` and `pthread_mutex_unlock` functions will fail if:

**EINVAL** The value specified by the `mutex` parameter does not refer to an initialized mutex object.

The `pthread_mutex_lock` function will fail if:

**EDEADLK** The current thread already owns the mutex and the mutex type is `PTHREAD_MUTEX_ERRORCHECK`.

The `pthread_mutex_unlock` function will fail if:

**EPERM** The current thread does not own the mutex and the mutex type is not `PTHREAD_MUTEX_NORMAL`.

These functions will not return an error code of `EINTR`.

## Related Information

The `pthread_mutex_init` or `pthread_mutex_destroy` (“`pthread_mutex_init` or `pthread_mutex_destroy` Subroutine” on page 1267) subroutine, `pthread.h` file.

Using Mutexes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_mutex\_timedlock Subroutine

### Purpose

Locks a mutex (ADVANCED REALTIME).

### Syntax

```
#include <pthread.h>
#include <time.h>

int pthread_mutex_timedlock(pthread_mutex_t *restrict mutex,
    const struct timespec *restrict abs_timeout);
```

### Description

The `pthread_mutex_timedlock()` function locks the mutex object referenced by `mutex`. If the mutex is already locked, the calling thread blocks until the mutex becomes available, as in the `pthread_mutex_lock()` function. If the mutex cannot be locked without waiting for another thread to unlock the mutex, this wait terminates when the specified timeout expires.

The timeout expires when the absolute time specified by `abs_timeout` passes—as measured by the clock on which timeouts are based (that is, when the value of that clock equals or exceeds `abs_timeout`)—or when the absolute time specified by `abs_timeout` has already been passed at the time of the call.

If the **Timers** option is supported, the timeout is based on the `CLOCK_REALTIME` clock; if the **Timers** option is not supported, the timeout is based on the system clock as returned by the `time()` function.



The resolution of the timeout matches the resolution of the clock on which it is based. The **timespec** data type is defined in the `<time.h>` header.

The function never fails with a timeout if the mutex can be locked immediately. The validity of the *abs\_timeout* parameter does not need to be checked if the mutex can be locked immediately.

As a consequence of the priority inheritance rules (for mutexes initialized with the `PRIO_INHERIT` protocol), if a timed mutex wait is terminated because its timeout expires, the priority of the owner of the mutex adjusts as necessary to reflect the fact that this thread is no longer among the threads waiting for the mutex.

## Application Usage

The `pthread_mutex_timedlock()` function is part of the **Threads** and **Timeouts** options and do not need to be provided on all implementations.

## Return Values

If successful, the `pthread_mutex_timedlock()` function returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_mutex_timedlock()` function fails if:

[EDEADLK]	The current thread already owns the mutex.
[EINVAL]	The mutex was created with the protocol attribute having the value <code>PTHREAD_PRIO_PROTECT</code> , and the calling thread's priority is higher than the mutex's current priority ceiling.
[EINVAL]	The process or thread would have blocked, and the <i>abs_timeout</i> parameter specified a nanoseconds field value less than 0 or greater than or equal to 1000 million.
[EINVAL]	<i>abs_timeout</i> is a NULL pointer.
[EINVAL]	The value specified by <i>mutex</i> does not refer to an initialized mutex object.
[ETIMEDOUT]	The mutex could not be locked before the specified timeout expired.

This function does not return an error code of [EINTR].

## Related Information

“mq\_receive, mq\_timedreceive Subroutine” on page 870, “posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event Subroutine” on page 1163, “pthread\_mutexattr\_destroy or pthread\_mutexattr\_init Subroutine,” “pthread\_mutex\_lock, pthread\_mutex\_trylock, or pthread\_mutex\_unlock Subroutine” on page 1270, “pthread\_rwlock\_timedrdlock Subroutine” on page 1287.

The `sem_timedwait` subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The `pthread.h` and `time.h` files in *AIX 5L Version 5.3 Files Reference*.

---

## pthread\_mutexattr\_destroy or pthread\_mutexattr\_init Subroutine

### Purpose

Initializes and destroys mutex attributes.

## Library

Threads Library (**libpthread.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_mutexattr_init (attr)
pthread_mutexattr_t *attr;
```

```
int pthread_mutexattr_destroy (attr)
pthread_mutexattr_t *attr;
```

## Description

The function **pthread\_mutexattr\_init** initializes a mutex attributes object *attr* with the default value for all of the attributes defined by the implementation.

The effect of initializing an already initialized mutex attributes object is undefined.

After a mutex attributes object has been used to initialize one or more mutexes, any function affecting the attributes object (including destruction) does not affect any previously initialized mutexes.

The **pthread\_mutexattr\_destroy** function destroys a mutex attributes object; the object becomes, in effect, uninitialized. An implementation may cause **pthread\_mutexattr\_destroy** to set the object referenced by *attr* to an invalid value. A destroyed mutex attributes object can be re-initialized using **pthread\_mutexattr\_init**; the results of otherwise referencing the object after it has been destroyed are undefined.

## Parameters

*attr* Specifies the mutex attributes object to initialize or delete.

## Return Values

Upon successful completion, **pthread\_mutexattr\_init** and **pthread\_mutexattr\_destroy** return zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_mutexattr\_init** function will fail if:

**ENOMEM** Insufficient memory exists to initialize the mutex attributes object.

The **pthread\_mutexattr\_destroy** function will fail if:

**EINVAL** The value specified by *attr* is invalid.

These functions will not return **EINTR**.

## Related Information

The **pthread\_create** (“pthread\_create Subroutine” on page 1243) subroutine, **pthread\_mutex\_init** or **pthread\_mutex\_destroy** (“pthread\_mutex\_init or pthread\_mutex\_destroy Subroutine” on page 1267) subroutine, **pthread\_cond\_destroy** or **pthread\_cond\_init** (“pthread\_cond\_destroy or pthread\_cond\_init Subroutine” on page 1233) subroutine, **pthread.h** file.

Using Mutexes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_mutexattr\_getkind\_np Subroutine

### Purpose

Returns the value of the kind attribute of a mutex attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_mutexattr_getkind_np ( attr, kind)
pthread_mutexattr_t *attr;
int *kind;
```

### Description

The **pthread\_mutexattr\_getkind\_np** subroutine returns the value of the kind attribute of the mutex attributes object *attr*. This attribute specifies the kind of the mutex created with this attributes object. It may have one of the following values:

<b>MUTEX_FAST_NP</b>	Denotes a fast mutex. A fast mutex can be locked only once. If the same thread unlocks twice the same fast mutex, the thread will deadlock. Any thread can unlock a fast mutex. A fast mutex is not compatible with the priority inheritance protocol.
<b>MUTEX_RECURSIVE_NP</b>	Denotes a recursive mutex. A recursive mutex can be locked more than once by the same thread without causing that thread to deadlock. The thread must then unlock the mutex as many times as it locked it. Only the thread that locked a recursive mutex can unlock it. A recursive mutex must not be used with condition variables.
<b>MUTEX_NONRECURSIVE_NP</b>	Denotes the default non-recursive POSIX compliant mutex.

### Notes:

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.
2. The **pthread\_mutexattr\_getkind\_np** subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

### Parameters

*attr* Specifies the mutex attributes object.  
*kind* Points to where the kind attribute value will be stored.

## Return Values

Upon successful completion, the value of the kind attribute is returned via the *kind* parameter, and 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_mutexattr_getkind_np` subroutine is unsuccessful if the following is true:

**EINVAL** The *attr* parameter is not valid.

## Related Information

The `pthread_mutexattr_setkind_np` (“`pthread_mutexattr_setkind_np` Subroutine” on page 1281) subroutine.

Using Mutexes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_mutexattr\_getprioceiling or pthread\_mutexattr\_setprioceiling Subroutine

### Purpose

Gets and sets the prioceiling attribute of the mutex attributes object.

### Syntax

```
#include <pthread.h>

int pthread_mutexattr_getprioceiling(const pthread_mutexattr_t *
    restrict attr, int *restrict prioceiling);
int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr,
    int prioceiling);
```

### Description

The `pthread_mutexattr_getprioceiling` and `pthread_mutexattr_setprioceiling` subroutines, respectively, get and set the priority ceiling attribute of a mutex attributes object pointed to by the *attr* parameter, which was previously created by the `pthread_mutexattr_init` subroutine.

The *prioceiling* attribute contains the priority ceiling of initialized mutexes. The values of the *prioceiling* parameter are within the maximum range of priorities defined by SCHED\_FIFO.

The *prioceiling* parameter defines the priority ceiling of initialized mutexes, which is the minimum priority level at which the critical section guarded by the mutex is executed. In order to avoid priority inversion, the priority ceiling of the mutex is set to a priority higher than or equal to the highest priority of all the threads that may lock that mutex. The values of the *prioceiling* parameter are within the maximum range of priorities defined under the SCHED\_FIFO scheduling policy.

### Return Values

Upon successful completion, the `pthread_mutexattr_getprioceiling` and `pthread_mutexattr_setprioceiling` subroutines return zero; otherwise, an error number shall be returned to indicate the error.

## Error Codes

The `pthread_mutexattr_getprioceiling` and `pthread_mutexattr_setprioceiling` subroutines can fail if:

<b>EINVAL</b>	The value specified by the <i>attr</i> or <i>prioceiling</i> parameter is invalid.
<b>ENOSYS</b>	This function is not supported (draft 7).
<b>ENOTSUP</b>	This function is not supported together with checkpoint/restart.
<b>EPERM</b>	The caller does not have the privilege to perform the operation in a strictly standards conforming environment where environment variable <b>XPG_SUS_ENV=ON</b> .

## Related Information

The “`pthread_mutex_init` or `pthread_mutex_destroy` Subroutine” on page 1267, “`pthread_mutex_lock`, `pthread_mutex_trylock`, or `pthread_mutex_unlock` Subroutine” on page 1270, “`pthread_mutex_timedlock` Subroutine” on page 1272.

The `pthread.h` file.

---

## `pthread_mutexattr_getprotocol` or `pthread_mutexattr_setprotocol` Subroutine

### Purpose

Gets and sets the protocol attribute of the mutex attributes object.

### Syntax

```
#include <pthread.h>

int pthread_mutexattr_getprotocol(const pthread_mutexattr_t *
    restrict attr, int *restrict protocol);
int pthread_mutexattr_setprotocol(pthread_mutexattr_t *attr,
    int protocol);
```

### Description

The `pthread_mutexattr_getprotocol` subroutine and `pthread_mutexattr_setprotocol` subroutine get and set the *protocol* parameter of a mutex attributes object pointed to by the *attr* parameter, which was previously created by the `pthread_mutexattr_init` subroutine.

The protocol attribute defines the protocol to be followed in utilizing mutexes. The value of the *protocol* parameter can be one of the following, which are defined in the `pthread.h` header file:

- **PTHREAD\_PRIO\_NONE**
- **PTHREAD\_PRIO\_INHERIT**
- **PTHREAD\_PRIO\_PROTECT**

When a thread owns a mutex with the **PTHREAD\_PRIO\_NONE** protocol attribute, its priority and scheduling are not affected by its mutex ownership.

When a thread is blocking higher priority threads because of owning one or more mutexes with the **PTHREAD\_PRIO\_INHERIT** protocol attribute, it executes at the higher of its priority or the priority of the highest priority thread waiting on any of the mutexes owned by this thread and initialized with this protocol.

When a thread owns one or more mutexes initialized with the **PTHREAD\_PRIO\_PROTECT** protocol, it executes at the higher of its priority or the highest of the priority ceilings of all the mutexes owned by this thread and initialized with this attribute, regardless of whether other threads are blocked on any of these mutexes. Privilege checking is necessary when the mutex priority ceiling is more favored than current

thread priority and the thread priority must be changed. The **pthread\_mutex\_lock** subroutine does not fail because of inappropriate privileges. Locking succeeds in this case, but no boosting is performed.

While a thread is holding a mutex which has been initialized with the **PTHREAD\_PRIO\_INHERIT** or **PTHREAD\_PRIO\_PROTECT** protocol attributes, it is not subject to being moved to the tail of the scheduling queue at its priority in the event that its original priority is changed, such as by a call to the **sched\_setparam** subroutine. Likewise, when a thread unlocks a mutex that has been initialized with the **PTHREAD\_PRIO\_INHERIT** or **PTHREAD\_PRIO\_PROTECT** protocol attributes, it is not subject to being moved to the tail of the scheduling queue at its priority in the event that its original priority is changed.

If a thread simultaneously owns several mutexes initialized with different protocols, it executes at the highest of the priorities that it would have obtained by each of these protocols.

When a thread makes a call to the **pthread\_mutex\_lock** subroutine, the mutex was initialized with the protocol attribute having the value **PTHREAD\_PRIO\_INHERIT**, when the calling thread is blocked because the mutex is owned by another thread, that owner thread inherits the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority and all its inherited priorities. Furthermore, if this owner thread itself becomes blocked on another mutex, the same priority inheritance effect shall be propagated to this other owner thread, in a recursive manner.

Behavior prior to AIX 5.3 is maintained under the non-POSIX protocol **PTHREAD\_PRIO\_DEFAULT**.

## Return Values

Upon successful completion, the **pthread\_mutexattr\_getprotocol** subroutine and the **pthread\_mutexattr\_setprotocol** subroutine return zero; otherwise, an error number shall be returned to indicate the error.

## Error Codes

The **pthread\_mutexattr\_setprotocol** subroutine fails if:

**ENOTSUP**            The value specified by the *protocol* parameter is an unsupported value.

The **pthread\_mutexattr\_getprotocol** subroutine and **pthread\_mutexattr\_setprotocol** subroutine can fail if:

**EINVAL**            The value specified by the *attr* parameter or the *protocol* parameter is invalid.  
**ENOSYS**            This function is not supported (draft 7).  
**ENOTSUP**            This function is not supported together with checkpoint/restart.  
**EPERM**             The caller does not have the privilege to perform the operation in a strictly standards conforming environment where environment variable **XPG\_SUS\_ENV=ON**.

## Related Information

The “**pthread\_mutex\_init** or **pthread\_mutex\_destroy** Subroutine” on page 1267, “**pthread\_mutex\_lock**, **pthread\_mutex\_trylock**, or **pthread\_mutex\_unlock** Subroutine” on page 1270, “**pthread\_mutex\_timedlock** Subroutine” on page 1272.

The **pthread.h** file.

---

## pthread\_mutexattr\_getpshared or pthread\_mutexattr\_setpshared Subroutine

### Purpose

Sets and gets process-shared attribute.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_mutexattr_getpshared (attr, pshared)
const pthread_mutexattr_t *attr;
int *pshared;
```

```
int pthread_mutexattr_setpshared (attr, pshared)
pthread_mutexattr_t *attr;
int pshared;
```

### Description

The **pthread\_mutexattr\_getpshared** subroutine obtains the value of the process-shared attribute from the attributes object referenced by *attr*. The **pthread\_mutexattr\_setpshared** subroutine is used to set the process-shared attribute in an initialized attributes object referenced by *attr*.

The process-shared attribute is set to `PTHREAD_PROCESS_SHARED` to permit a mutex to be operated upon by any thread that has access to the memory where the mutex is allocated, even if the mutex is allocated in memory that is shared by multiple processes. If the **process-shared** attribute is `PTHREAD_PROCESS_PRIVATE`, the mutex will only be operated upon by threads created within the same process as the thread that initialized the mutex; if threads of differing processes attempt to operate on such a mutex, the behavior is undefined. The default value of the attribute is `PTHREAD_PROCESS_PRIVATE`.

### Parameters

*attr*                Specifies the mutex attributes object.  
*pshared*            Points to where the pshared attribute value will be stored.

### Return Values

Upon successful completion, the **pthread\_mutexattr\_setpshared** subroutine returns zero. Otherwise, an error number is returned to indicate the error.

Upon successful completion, the **pthread\_mutexattr\_getpshared** subroutine returns zero and stores the value of the process-shared attribute of *attr* into the object referenced by the *pshared* parameter. Otherwise, an error number is returned to indicate the error.

### Error Codes

The **pthread\_mutexattr\_getpshared** and **pthread\_mutexattr\_setpshared** subroutines will fail if:

**EINVAL**        The value specified by *attr* is invalid.

The `pthread_mutexattr_setpshared` function will fail if:

**EINVAL** The new value specified for the attribute is outside the range of legal values for that attribute.

These subroutines will not return an error code of `EINTR`.

## Related Information

The `pthread_mutexattr_init` (“`pthread_mutexattr_destroy` or `pthread_mutexattr_init` Subroutine” on page 1273) subroutine.

Advanced Attributes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_mutexattr\_gettype or pthread\_mutexattr\_settype Subroutine

### Purpose

Gets or sets a mutex type.

### Library

Threads Library (`libthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_mutexattr_gettype (attr, type)
const pthread_mutexattr_t *attr;
int *type;
```

```
int pthread_mutexattr_settype (attr, type)
pthread_mutexattr_t *attr;
int type;
```

### Description

The `pthread_mutexattr_gettype` and `pthread_mutexattr_settype` subroutines respectively get and set the mutex type attribute. This attribute is set in the `type` parameter to these subroutines. The default value of the type attribute is `PTHREAD_MUTEX_DEFAULT`. The type of mutex is contained in the type attribute of the mutex attributes. Valid mutex types include:

#### **PTHREAD\_MUTEX\_NORMAL**

This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.

#### **PTHREAD\_MUTEX\_ERRORCHECK**

This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex which another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.



## PTHREAD\_MUTEX\_RECURSIVE

A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock which can occur with mutexes of type `PTHREAD_MUTEX_NORMAL` cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex which another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.

## PTHREAD\_MUTEX\_DEFAULT

Attempting to recursively lock a mutex of this type results in undefined behavior. Attempting to unlock a mutex of this type which was not locked by the calling thread results in undefined behavior. Attempting to unlock a mutex of this type which is not locked results in undefined behavior. An implementation is allowed to map this mutex to one of the other mutex types.

It is advised that an application should not use a `PTHREAD_MUTEX_RECURSIVE` mutex with condition variables because the implicit unlock performed for a `pthread_cond_wait` or `pthread_cond_timedwait` may not actually release the mutex (if it had been locked multiple times). If this happens, no other thread can satisfy the condition of the predicate.

## Parameters

*attr* Specifies the mutex object to get or set.  
*type* Specifies the type to get or set.

## Return Values

If successful, the `pthread_mutexattr_settype` subroutine returns zero. Otherwise, an error number is returned to indicate the error. Upon successful completion, the `pthread_mutexattr_gettype` subroutine returns zero and stores the value of the `type` attribute of *attr* into the object referenced by the *type* parameter. Otherwise an error is returned to indicate the error.

## Error Codes

The `pthread_mutexattr_gettype` and `pthread_mutexattr_settype` subroutines will fail if:

**EINVAL** The value of the *type* parameter is invalid.  
**EINVAL** The value specified by the *attr* parameter is invalid.

## Related Information

The `pthread_cond_wait` (“`pthread_cond_wait` or `pthread_cond_timedwait` Subroutine” on page 1236) and `pthread_cond_timedwait` (“`pthread_cond_wait` or `pthread_cond_timedwait` Subroutine” on page 1236) subroutines.

The `pthread.h` file.

---

## pthread\_mutexattr\_setkind\_np Subroutine

### Purpose

Sets the value of the `kind` attribute of a mutex attributes object.

## Library

Threads Library (**libpthreads.a**)

## Syntax

```
#include <pthread.h>
```

```
int pthread_mutexattr_setkind_np ( attr, kind)
pthread_mutexattr_t *attr;
int kind;
```

## Description

The **pthread\_mutexattr\_setkind\_np** subroutine sets the value of the *kind* attribute of the mutex attributes object *attr*. This attribute specifies the kind of the mutex created with this attributes object.

### Notes:

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.
2. The **pthread\_mutexattr\_setkind\_np** subroutine is not portable.

This subroutine is provided only for compatibility with the DCE threads. It should not be used when writing new applications.

## Parameters

*attr* Specifies the mutex attributes object.  
*kind* Specifies the kind to set. It must have one of the following values:

### **MUTEX\_FAST\_NP**

Denotes a fast mutex. A fast mutex can be locked only once. If the same thread unlocks twice the same fast mutex, the thread will deadlock. Any thread can unlock a fast mutex. A fast mutex is not compatible with the priority inheritance protocol.

### **MUTEX\_RECURSIVE\_NP**

Denotes a recursive mutex. A recursive mutex can be locked more than once by the same thread without causing that thread to deadlock. The thread must then unlock the mutex as many times as it locked it. Only the thread that locked a recursive mutex can unlock it. A recursive mutex must not be used with condition variables.

### **MUTEX\_NONRECURSIVE\_NP**

Denotes the default non-recursive POSIX compliant mutex.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The **pthread\_mutexattr\_setkind\_np** subroutine is unsuccessful if the following is true:

**EINVAL** The *attr* parameter is not valid.  
**ENOTSUP** The value of the *kind* parameter is not supported.

## Related Information

The **pthread\_mutexattr\_getkind\_np** (“**pthread\_mutexattr\_getkind\_np** Subroutine” on page 1275) subroutine.

## pthread\_once Subroutine

### Purpose

Executes a routine exactly once in a process.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>

int pthread_once (once_control, init_routine)
pthread_once_t *once_control;
void (*init_routine)(void);
,
pthread_once_t once_control = PTHREAD_ONCE_INIT;
```

### Description

The **pthread\_once** subroutine executes the routine *init\_routine* exactly once in a process. The first call to this subroutine by any thread in the process executes the given routine, without parameters. Any subsequent call will have no effect.

The *init\_routine* routine is typically an initialization routine. Multiple initializations can be handled by multiple instances of **pthread\_once\_t** structures. This subroutine is useful when a unique initialization has to be done by one thread among many. It reduces synchronization requirements.

**Note:** The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

### Parameters

<i>once_control</i>	Points to a synchronization control structure. This structure has to be initialized by the static initializer macro <b>PTHREAD_ONCE_INIT</b> .
<i>init_routine</i>	Points to the routine to be executed.

### Return Values

Upon successful completion, **pthread\_once** returns zero. Otherwise, an error number is returned to indicate the error.

### Error Codes

No errors are defined. The **pthread\_once** function will not return an error code of **EINTR**.

### Related Information

The **pthread\_create** (“pthread\_create Subroutine” on page 1243) subroutine, **pthread.h** file, **PTHREAD\_ONCE\_INIT** (“PTHREAD\_ONCE\_INIT Macro” on page 1284) macro.

One Time Initializations in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## PTHREAD\_ONCE\_INIT Macro

### Purpose

Initializes a once synchronization control structure.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
static pthread_once_t once_block = PTHREAD_ONCE_INIT;
```

### Description

The **PTHREAD\_ONCE\_INIT** macro initializes the static once synchronization control structure *once\_block*, used for one-time initializations with the **pthread\_once** (“pthread\_once Subroutine” on page 1283) subroutine. The once synchronization control structure must be static to ensure the unicity of the initialization.

**Note:** The **pthread.h** file header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the **cc\_r** compiler used. In this case, the flag is automatically set.

### Related Information

The **pthread\_once** (“pthread\_once Subroutine” on page 1283) subroutine.

One Time Initializations in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_rwlock\_init or pthread\_rwlock\_destroy Subroutine

### Purpose

Initializes or destroys a read-write lock object.

### Library

Threads Library (**libthreads.a**)

### Syntax

```
#include <pthread.h>

int pthread_rwlock_init (rwlock, attr)
pthread_rwlock_t *rwlock;
const pthread_rwlockattr_t *attr;

int pthread_rwlock_destroy (rwlock)
pthread_rwlock_t *rwlock;
pthread_rwlock_t rwlock=PTHREAD_RWLOCK_INITIALIZER;
```

### Description

The **pthread\_rwlock\_init** subroutine initializes the read-write lock referenced by *rwlock* with the attributes referenced by *attr*. If *attr* is NULL, the default read-write lock attributes are used; the effect is the same as passing the address of a default read-write lock attributes object. Once initialized, the lock can be used

any number of times without being re-initialized. Upon successful initialization, the state of the read-write lock becomes initialized and unlocked. Results are undefined if **pthread\_rwlock\_init** is called specifying an already initialized read-write lock. Results are undefined if a read-write lock is used without first being initialized.

If the **pthread\_rwlock\_init** function fails, *rwlock* is not initialized and the contents of *rwlock* are undefined.

The **pthread\_rwlock\_destroy** function destroys the read-write lock object referenced by *rwlock* and releases any resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is re-initialized by another call to **pthread\_rwlock\_init**. An implementation may cause **pthread\_rwlock\_destroy** to set the object referenced by *rwlock* to an invalid value. Results are undefined if **pthread\_rwlock\_destroy** is called when any thread holds *rwlock*. Attempting to destroy an uninitialized read-write lock results in undefined behavior. A destroyed read-write lock object can be re-initialized using **pthread\_rwlock\_init**; the results of otherwise referencing the read-write lock object after it has been destroyed are undefined.

In cases where default read-write lock attributes are appropriate, the macro **PTHREAD\_RWLOCK\_INITIALIZER** can be used to initialize read-write locks that are statically allocated. The effect is equivalent to dynamic initialization by a call to **pthread\_rwlock\_init** with the parameter *attr* specified as NULL, except that no error checks are performed.

## Parameters

*rwlock* Specifies the read-write lock to be initialized or destroyed.  
*attr* Specifies the attributes of the read-write lock to be initialized.

## Return Values

If successful, the **pthread\_rwlock\_init** and **pthread\_rwlock\_destroy** functions return zero. Otherwise, an error number is returned to indicate the error. The EBUSY and EINVAL error checks, if implemented, will act as if they were performed immediately at the beginning of processing for the function and caused an error return prior to modifying the state of the read-write lock specified by *rwlock*.

## Error Codes

The **pthread\_rwlock\_init** subroutine will fail if:

**ENOMEM** Insufficient memory exists to initialize the read-write lock.  
**EINVAL** The value specified by *attr* is invalid.

The **pthread\_rwlock\_destroy** subroutine will fail if:

**EBUSY** The implementation has detected an attempt to destroy the object referenced by *rwlock* while it is locked.  
**EINVAL** The value specified by *attr* is invalid.

## Related Information

The **pthread.h** file.

The **pthread\_rwlock\_rdlock** (“pthread\_rwlock\_rdlock or pthread\_rwlock\_tryrdlock Subroutines” on page 1286), **pthread\_rwlock\_wrlock** (“pthread\_rwlock\_wrlock or pthread\_rwlock\_trywrlock Subroutines” on page 1291), **pthread\_rwlockattr\_init** (“pthread\_rwlockattr\_init or pthread\_rwlockattr\_destroy Subroutines” on page 1293) and **pthread\_rwlock\_unlock** (“pthread\_rwlock\_unlock Subroutine” on page 1290) subroutines.

---

## pthread\_rwlock\_rdlock or pthread\_rwlock\_tryrdlock Subroutines

### Purpose

Locks a read-write lock object for reading.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_rwlock_rdlock (rwlock)
pthread_rwlock_t *rwlock;
```

```
int pthread_rwlock_tryrdlock (rwlock)
pthread_rwlock_t *rwlock;
```

### Description

The **pthread\_rwlock\_rdlock** function applies a read lock to the read-write lock referenced by *rwlock*. The calling thread acquires the read lock if a writer does not hold the lock and there are no writers blocked on the lock. It is unspecified whether the calling thread acquires the lock when a writer does not hold the lock and there are writers waiting for the lock. If a writer holds the lock, the calling thread will not acquire the read lock. If the read lock is not acquired, the calling thread blocks (that is, it does not return from the **pthread\_rwlock\_rdlock call**) until it can acquire the lock. Results are undefined if the calling thread holds a write lock on *rwlock* at the time the call is made.

Implementations are allowed to favor writers over readers to avoid writer starvation.

A thread may hold multiple concurrent read locks on *rwlock* (that is, successfully call the **pthread\_rwlock\_rdlock** function *n* times). If so, the thread must perform matching unlocks (that is, it must call the **pthread\_rwlock\_unlock** function *n* times).

The function **pthread\_rwlock\_tryrdlock** applies a read lock as in the **pthread\_rwlock\_rdlock** function with the exception that the function fails if any thread holds a write lock on *rwlock* or there are writers blocked on *rwlock*.

Results are undefined if any of these functions are called with an uninitialized read-write lock.

If a signal is delivered to a thread waiting for a read-write lock for reading, upon return from the signal handler the thread resumes waiting for the read-write lock for reading as if it was not interrupted.

### Parameters

*rwlock*                                 Specifies the read-write lock to be locked for reading.

### Return Values

If successful, the **pthread\_rwlock\_rdlock** function returns zero. Otherwise, an error number is returned to indicate the error.

The function **pthread\_rwlock\_tryrdlock** returns zero if the lock for reading on the read-write lock object referenced by *rwlock* is acquired. Otherwise an error number is returned to indicate the error.

## Error Codes

The `pthread_rwlock_tryrdlock` function will fail if:

**EBUSY** The read-write lock could not be acquired for reading because a writer holds the lock or was blocked on it.

The `pthread_rwlock_rdlock` and `pthread_rwlock_tryrdlock` functions will fail if:

**EINVAL** The value specified by *rwlock* does not refer to an initialized read-write lock object.

**EDEADLK** The current thread already owns the read-write lock for writing.

**EAGAIN** The read lock could not be acquired because the maximum number of read locks for *rwlock* has been exceeded.

## Implementation Specifics

Realtime applications may encounter priority inversion when using read-write locks. The problem occurs when a high priority thread 'locks' a read-write lock that is about to be 'unlocked' by a low priority thread, but the low priority thread is preempted by a medium priority thread. This scenario leads to priority inversion; a high priority thread is blocked by lower priority threads for an unlimited period of time. During system design, realtime programmers must take into account the possibility of this kind of priority inversion. They can deal with it in a number of ways, such as by having critical sections that are guarded by read-write locks execute at a high priority, so that a thread cannot be preempted while executing in its critical section.

## Related Information

The `pthread.h` file.

The `pthread_rwlock_init` (“`pthread_rwlock_init` or `pthread_rwlock_destroy` Subroutine” on page 1284), `pthread_rwlock_wrlock` (“`pthread_rwlock_wrlock` or `pthread_rwlock_trywrlock` Subroutines” on page 1291), `pthread_rwlockattr_init` (“`pthread_rwlockattr_init` or `pthread_rwlockattr_destroy` Subroutines” on page 1293), and `pthread_rwlock_unlock` (“`pthread_rwlock_unlock` Subroutine” on page 1290) subroutines.

---

## pthread\_rwlock\_timedrdlock Subroutine

### Purpose

Locks a read-write lock for reading.

### Syntax

```
#include <pthread.h>
#include <time.h>
```

```
int pthread_rwlock_timedrdlock(pthread_rwlock_t *restrict rwlock,
                               const struct timespec *restrict abs_timeout);
```

### Description

The `pthread_rwlock_timedrdlock()` function applies a read lock to the read-write lock referenced by *rwlock* as in the `pthread_rwlock_rdlock()` function. However, if the lock cannot be acquired without waiting for other threads to unlock the lock, this wait terminates when the specified timeout expires. The timeout expires when the absolute time specified by *abs\_timeout* passes—as measured by the clock on which timeouts are based (that is, when the value of that clock equals or exceeds *abs\_timeout*)—or when the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

If the **Timers** option is supported, the timeout is based on the `CLOCK_REALTIME` clock; if the **Timers** option is not supported, the timeout is based on the system clock as returned by the `time()` function.

The resolution of the timeout matches the resolution of the clock on which it is based. The `timespec` data type is defined in the `<time.h>` header.

The function never fails with a timeout if the lock can be acquired immediately. The validity of the `abs_timeout` parameter does not need to be checked if the lock can be immediately acquired.

If a signal that causes a signal handler to be executed is delivered to a thread that is blocked on a read-write lock through a call to `pthread_rwlock_timedrdlock()`, the thread resumes waiting for the lock (as if it were not interrupted) after the signal handler returns.

The calling thread can deadlock if it holds a write lock on **rwlock** at the time the call is made. The results are undefined if this function is called with an uninitialized read-write lock.

## Application Usage

The `pthread_rwlock_timedrdlock()` function is part of the **Threads** and **Timeouts** options and do not need to be provided on all implementations.

## Return Values

The `pthread_rwlock_timedrdlock()` function returns 0 if the lock for reading on the read-write lock object referenced by `rwlock` is acquired. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_rwlock_timedrdlock()` function fails if:

[ETIMEDOUT]                      The lock could not be acquired before the specified timeout expired.

The `pthread_rwlock_timedrdlock()` function might fail if:

[EAGAIN]                         The read lock could not be acquired because the maximum number of read locks for lock would be exceeded.

[EDEADLK]                        The calling thread already holds a write lock on `rwlock`.

[EINVAL]                         The value specified by `rwlock` does not refer to an initialized read-write lock object, or the `abs_timeout` nanosecond value is less than 0 or greater than or equal to 1000 million.

This function does not return an error code of [EINTR].

## Related Information

“mq\_receive, mq\_timedreceive Subroutine” on page 870, “posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event Subroutine” on page 1163, “pthread\_mutex\_timedlock Subroutine” on page 1272, “pthread\_rwlock\_init or pthread\_rwlock\_destroy Subroutine” on page 1284, “pthread\_rwlock\_rdlock or pthread\_rwlock\_tryrdlock Subroutines” on page 1286, “pthread\_rwlock\_timedwrlock Subroutine” on page 1289, “pthread\_rwlock\_wrlock or pthread\_rwlock\_trywrlock Subroutines” on page 1291, “pthread\_rwlock\_unlock Subroutine” on page 1290.

The `sem_timedwait` subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The `pthread.h` and `time.h` files in *AIX 5L Version 5.3 Files Reference*.



---

## pthread\_rwlock\_timedwrlock Subroutine

### Purpose

Locks a read-write lock for writing.

### Syntax

```
#include <pthread.h>
#include <time.h>

int pthread_rwlock_timedwrlock(pthread_rwlock_t *restrict rwlock,
                               const struct timespec *restrict abs_timeout);
```

### Description

The **pthread\_rwlock\_timedwrlock()** function applies a write lock to the read-write lock referenced by *rwlock* as in the **pthread\_rwlock\_wrlock()** function. However, if the lock cannot be acquired without waiting for other threads to unlock the lock, this wait terminates when the specified timeout expires. The timeout expires when the absolute time specified by *abs\_timeout* passes—as measured by the clock on which timeouts are based (that is, when the value of that clock equals or exceeds *abs\_timeout*)—or when the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

If the **Timers** option is supported, the timeout is based on the CLOCK\_REALTIME clock; if the **Timers** option is not supported, the timeout is based on the system clock as returned by the **time()** function.

The resolution of the timeout matches the resolution of the clock on which it is based. The **timespec** data type is defined in the **<time.h>** header.

The function never fails with a timeout if the lock can be acquired immediately. The validity of the *abs\_timeout* parameter does not need to be checked if the lock can be immediately acquired.

If a signal that causes a signal handler to be executed is delivered to a thread that is blocked on a read-write lock through a call to **pthread\_rwlock\_timedwrlock()**, the thread resumes waiting for the lock (as if it were not interrupted) after the signal handler returns.

The calling thread can deadlock if it holds the read-write lock at the time the call is made. The results are undefined if this function is called with an uninitialized read-write lock.

### Application Usage

The **pthread\_rwlock\_timedwrlock()** function is part of the **Threads** and **Timeouts** options and do not need to be provided on all implementations.

### Return Values

The **pthread\_rwlock\_timedwrlock()** function returns 0 if the lock for writing on the read-write lock object referenced by *rwlock* is acquired. Otherwise, an error number is returned to indicate the error.

### Error Codes

The **pthread\_rwlock\_timedrdlock()** function fails if:

**ETIMEDOUT**                      The lock could not be acquired before the specified timeout expired.

The **pthread\_rwlock\_timedrdlock()** function might fail if:

**EDEADLK**                        The calling thread already holds the *rwlock*.

**EINVAL** The value specified by *rwlock* does not refer to an initialized read-write lock object, or the *abs\_timeout* nanosecond value is less than 0 or greater than or equal to 1000 million.

This function does not return an error code of **EINTR**.

## Related Information

“mq\_receive, mq\_timedreceive Subroutine” on page 870, “posix\_trace\_getnext\_event, posix\_trace\_timedgetnext\_event, posix\_trace\_trygetnext\_event Subroutine” on page 1163, “pthread\_mutex\_timedlock Subroutine” on page 1272, “pthread\_rwlock\_init or pthread\_rwlock\_destroy Subroutine” on page 1284, “pthread\_rwlock\_rdlock or pthread\_rwlock\_tryrdlock Subroutines” on page 1286, “pthread\_rwlock\_wrlock or pthread\_rwlock\_trywrlock Subroutines” on page 1291, “pthread\_rwlock\_unlock Subroutine.”

The **sem\_timedwait** subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The **pthread.h** and **time.h** files in *AIX 5L Version 5.3 Files Reference*.

---

## pthread\_rwlock\_unlock Subroutine

### Purpose

Unlocks a read-write lock object.

### Library

Threads Library (**libthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_rwlock_unlock (rwlock)  
pthread_rwlock_t *rwlock;
```

### Description

The **pthread\_rwlock\_unlock** subroutine is called to release a lock held on the read-write lock object referenced by *rwlock*. Results are undefined if the read-write lock *rwlock* is not held by the calling thread.

If this subroutine is called to release a read lock from the read-write lock object and there are other read locks currently held on this read-write lock object, the read-write lock object remains in the read locked state. If this subroutine releases the calling thread's last read lock on this read-write lock object, then the calling thread is no longer one of the owners of the object. If this subroutine releases the last read lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.

If this subroutine is called to release a write lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.

If the call to the **pthread\_rwlock\_unlock** subroutine results in the read-write lock object becoming unlocked and there are multiple threads waiting to acquire the read-write lock object for writing, the scheduling policy is used to determine which thread acquires the read-write lock object for writing. If there are multiple threads waiting to acquire the read-write lock object for reading, the scheduling policy is used to determine the order in which the waiting threads acquire the read-write lock object for reading. If there

are multiple threads blocked on *rwlock* for both read locks and write locks, it is unspecified whether the readers acquire the lock first or whether a writer acquires the lock first.

Results are undefined if any of these subroutines are called with an uninitialized read-write lock.

## Parameters

*rwlock* Specifies the read-write lock to be unlocked.

## Return Values

If successful, the **pthread\_rwlock\_unlock** subroutine returns zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The **pthread\_rwlock\_unlock** subroutine may fail if:

**EINVAL** The value specified by *rwlock* does not refer to an initialized read-write lock object.  
**EPERM** The current thread does not own the read-write lock.

## Related Information

The **pthread.h** file.

The **pthread\_rwlock\_init** (“**pthread\_rwlock\_init** or **pthread\_rwlock\_destroy** Subroutine” on page 1284), **pthread\_rwlock\_wrlock** (“**pthread\_rwlock\_wrlock** or **pthread\_rwlock\_trywrlock** Subroutines”), **pthread\_rwlockattr\_init** (“**pthread\_rwlockattr\_init** or **pthread\_rwlockattr\_destroy** Subroutines” on page 1293), **pthread\_rwlock\_rdlock** (“**pthread\_rwlock\_rdlock** or **pthread\_rwlock\_tryrdlock** Subroutines” on page 1286) subroutines.

---

## pthread\_rwlock\_wrlock or pthread\_rwlock\_trywrlock Subroutines

### Purpose

Locks a read-write lock object for writing.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_rwlock_wrlock (rwlock)  
pthread_rwlock_t *rwlock;
```

```
int pthread_rwlock_trywrlock (rwlock)  
pthread_rwlock_t *rwlock;
```

### Description

The **pthread\_rwlock\_wrlock** subroutine applies a write lock to the read-write lock referenced by *rwlock*. The calling thread acquires the write lock if no other thread (reader or writer) holds the read-write lock *rwlock*. Otherwise, the thread blocks (that is, does not return from the **pthread\_rwlock\_wrlock** call) until it can acquire the lock. Results are undefined if the calling thread holds the read-write lock (whether a read or write lock) at the time the call is made.

Implementations are allowed to favor writers over readers to avoid writer starvation.

The **pthread\_rwlock\_trywrlock** subroutine applies a write lock like the **pthread\_rwlock\_wrlock** subroutine, with the exception that the function fails if any thread currently holds *rwlock* (for reading or writing).

Results are undefined if any of these functions are called with an uninitialized read-write lock.

If a signal is delivered to a thread waiting for a read-write lock for writing, upon return from the signal handler the thread resumes waiting for the read-write lock for writing as if it was not interrupted.

Realtime applications may encounter priority inversion when using read-write locks. The problem occurs when a high priority thread 'locks' a read-write lock that is about to be 'unlocked' by a low priority thread, but the low priority thread is preempted by a medium priority thread. This scenario leads to priority inversion; a high priority thread is blocked by lower priority threads for an unlimited period of time. During system design, realtime programmers must take into account the possibility of this kind of priority inversion. They can deal with it in a number of ways, such as by having critical sections that are guarded by read-write locks execute at a high priority, so that a thread cannot be preempted while executing in its critical section.

## Parameters

*rwlock* Specifies the read-write lock to be locked for writing.

## Return Values

If successful, the **pthread\_rwlock\_wrlock** subroutine returns zero. Otherwise, an error number is returned to indicate the error.

The **pthread\_rwlock\_trywrlock** subroutine returns zero if the lock for writing on the read-write lock object referenced by *rwlock* is acquired. Otherwise an error number is returned to indicate the error.

## Error Codes

The **pthread\_rwlock\_trywrlock** subroutine will fail if:

**EBUSY** The read-write lock could not be acquired for writing because it was already locked for reading or writing.

The **pthread\_rwlock\_wrlock** and **pthread\_rwlock\_trywrlock** subroutines may fail if:

**EINVAL** The value specified by *rwlock* does not refer to an initialized read-write lock object.

**EDEADLK** The current thread already owns the read-write lock for writing or reading.

## Related Information

The **pthread.h** file.

The **pthread\_rwlock\_init** ("pthread\_rwlock\_init or pthread\_rwlock\_destroy Subroutine" on page 1284), **pthread\_rwlock\_unlock** ("pthread\_rwlock\_unlock Subroutine" on page 1290), **pthread\_rwlockattr\_init** ("pthread\_rwlockattr\_init or pthread\_rwlockattr\_destroy Subroutines" on page 1293), **pthread\_rwlock\_rdlock** ("pthread\_rwlock\_rdlock or pthread\_rwlock\_tryrdlock Subroutines" on page 1286) subroutines.

---

## pthread\_rwlockattr\_init or pthread\_rwlockattr\_destroy Subroutines

### Purpose

Initializes and destroys read-write lock attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_rwlockattr_init (attr)
pthread_rwlockattr_t *attr;
```

```
int pthread_rwlockattr_destroy (attr)
pthread_rwlockattr_t *attr;
```

### Description

The **pthread\_rwlockattr\_init** subroutine initializes a read-write lock attributes object *attr* with the default value for all of the attributes defined by the implementation. Results are undefined if **pthread\_rwlockattr\_init** is called specifying an already initialized read-write lock attributes object.

After a read-write lock attributes object has been used to initialize one or more read-write locks, any function affecting the attributes object (including destruction) does not affect any previously initialized read-write locks.

The **pthread\_rwlockattr\_destroy** subroutine destroys a read-write lock attributes object. The effect of subsequent use of the object is undefined until the object is re-initialized by another call to **pthread\_rwlockattr\_init**. An implementation may cause **pthread\_rwlockattr\_destroy** to set the object referenced by *attr* to an invalid value.

### Parameters

*attr*                                Specifies a read-write lock attributes object to be initialized or destroyed.

### Return Value

If successful, the **pthread\_rwlockattr\_init** and **pthread\_rwlockattr\_destroy** subroutines return zero. Otherwise, an error number is returned to indicate the error.

### Error Codes

The **pthread\_rwlockattr\_init** subroutine will fail if:

**ENOMEM**    Insufficient memory exists to initialize the read-write lock attributes object.

The **pthread\_rwlockattr\_destroy** subroutine will fail if:

**EINVAL**     The value specified by *attr* is invalid.

### Related Information

The **pthread.h** file.

The **pthread\_rwlock\_init** (“pthread\_rwlock\_init or pthread\_rwlock\_destroy Subroutine” on page 1284), **pthread\_rwlock\_unlock** (“pthread\_rwlock\_unlock Subroutine” on page 1290), **pthread\_rwlock\_wrlock** (“pthread\_rwlock\_wrlock or pthread\_rwlock\_trywrlock Subroutines” on page 1291), **pthread\_rwlock\_rdlock** (“pthread\_rwlock\_rdlock or pthread\_rwlock\_tryrdlock Subroutines” on page 1286), and **pthread\_rwlockattr\_getpshared** (“pthread\_rwlockattr\_getpshared or pthread\_rwlockattr\_setpshared Subroutines”) subroutines.

---

## pthread\_rwlockattr\_getpshared or pthread\_rwlockattr\_setpshared Subroutines

### Purpose

Gets and sets process-shared attribute of read-write lock attributes object.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
int pthread_rwlockattr_getpshared (attr, pshared)
const pthread_rwlockattr_t *attr;
int *pshared;
```

```
int pthread_rwlockattr_setpshared (attr, pshared)
pthread_rwlockattr_t *attr;
int pshared;
```

### Description

The process-shared attribute is set to `PTHREAD_PROCESS_SHARED` to permit a read-write lock to be operated upon by any thread that has access to the memory where the read-write lock is allocated, even if the read-write lock is allocated in memory that is shared by multiple processes. If the process-shared attribute is `PTHREAD_PROCESS_PRIVATE`, the read-write lock will only be operated upon by threads created within the same process as the thread that initialized the read-write lock; if threads of differing processes attempt to operate on such a read-write lock, the behavior is undefined. The default value of the process-shared attribute is `PTHREAD_PROCESS_PRIVATE`.

The **pthread\_rwlockattr\_getpshared** subroutine obtains the value of the process-shared attribute from the initialized attributes object referenced by *attr*. The **pthread\_rwlockattr\_setpshared** subroutine is used to set the process-shared attribute in an initialized attributes object referenced by *attr*.

### Parameters

<i>attr</i>	Specifies the initialized attributes object.
<i>pshared</i>	Specifies the process-shared attribute of read-write lock attributes object to be gotten and set.

### Return Values

If successful, the **pthread\_rwlockattr\_setpshared** subroutine returns zero. Otherwise, an error number is returned to indicate the error.

Upon successful completion, the **pthread\_rwlockattr\_getpshared** subroutine returns zero and stores the value of the process-shared attribute of *attr* into the object referenced by the *pshared* parameter. Otherwise an error number is returned to indicate the error.

## Error Codes

The `pthread_rwlockattr_getpshared` and `pthread_rwlockattr_setpshared` subroutines will fail if:

**EINVAL** The value specified by *attr* is invalid.

The `pthread_rwlockattr_setpshared` subroutine will fail if:

**EINVAL** The new value specified for the attribute is outside the range of legal values for that attribute.

## Related Information

The `pthread.h` file.

The `pthread_rwlock_init` (“`pthread_rwlock_init` or `pthread_rwlock_destroy` Subroutine” on page 1284), `pthread_rwlock_unlock` (“`pthread_rwlock_unlock` Subroutine” on page 1290), `pthread_rwlock_wrlock` (“`pthread_rwlock_wrlock` or `pthread_rwlock_trywrlock` Subroutines” on page 1291), `pthread_rwlock_rdlock` (“`pthread_rwlock_rdlock` or `pthread_rwlock_tryrdlock` Subroutines” on page 1286), `pthread_rwlockattr_init` (“`pthread_rwlockattr_init` or `pthread_rwlockattr_destroy` Subroutines” on page 1293) subroutines.

---

## pthread\_self Subroutine

### Purpose

Returns the calling thread’s ID.

### Library

Threads Library (`libpthread.a`)

### Syntax

```
#include <pthread.h>
pthread_t pthread_self (void);
```

### Description

The `pthread_self` subroutine returns the calling thread’s ID.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

### Return Values

The calling thread’s ID is returned.

### Errors

No errors are defined.

The `pthread_self` function will not return an error code of `EINTR`.

## Related Information

The `pthread_create` (“`pthread_create` Subroutine” on page 1243) subroutine, `pthread_equal` (“`pthread_equal` Subroutine” on page 1247) subroutine.

The `pthread.h` file.

Creating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_setcancelstate, pthread\_setcanceltype, or pthread\_testcancel Subroutines

### Purpose

Sets the calling thread's cancelability state.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_setcancelstate (state, oldstate)
int state;
int *oldstate;
```

```
int pthread_setcanceltype (type, oldtype)
int type;
int *oldtype;
```

```
int pthread_testcancel (void)
```

### Description

The `pthread_setcancelstate` subroutine atomically both sets the calling thread's cancelability state to the indicated state and returns the previous cancelability state at the location referenced by `oldstate`. Legal values for state are `PTHREAD_CANCEL_ENABLE` and `PTHREAD_CANCEL_DISABLE`.

The `pthread_setcanceltype` subroutine atomically both sets the calling thread's cancelability type to the indicated type and returns the previous cancelability type at the location referenced by `oldtype`. Legal values for type are `PTHREAD_CANCEL_DEFERRED` and `PTHREAD_CANCEL_ASYNCHEOUS`.

The cancelability state and type of any newly created threads, including the thread in which `main` was first invoked, are `PTHREAD_CANCEL_ENABLE` and `PTHREAD_CANCEL_DEFERRED` respectively.

The `pthread_testcancel` subroutine creates a cancellation point in the calling thread. The `pthread_testcancel` subroutine has no effect if cancelability is disabled.

### Parameters

<i>state</i>	Specifies the new cancelability state to set. It must have one of the following values: <b>PTHREAD_CANCEL_DISABLE</b> Disables cancelability; the thread is not cancelable. Cancellation requests are held pending. <b>PTHREAD_CANCEL_ENABLE</b> Enables cancelability; the thread is cancelable, according to its cancelability type. This is the default value.
<i>oldstate</i>	Points to where the previous cancelability state value will be stored.
<i>type</i>	Specifies the new cancelability type to set.
<i>oldtype</i>	Points to where the previous cancelability type value will be stored.



## Return Values

If successful, the `pthread_setcancelstate` and `pthread_setcanceltype` subroutines return zero. Otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_setcancelstate` subroutine will fail if:

**EINVAL** The specified state is not `PTHREAD_CANCEL_ENABLE` or `PTHREAD_CANCEL_DISABLE`.

The `pthread_setcanceltype` subroutine will fail if:

**EINVAL** The specified type is not `PTHREAD_CANCEL_DEFERRED` or `PTHREAD_CANCEL_ASYNCHRONOUS`.

These subroutines will not return an error code of `EINTR`.

## Related Information

The `pthread_cancel` (“`pthread_cancel` Subroutine” on page 1231) subroutine.

The `pthread.h` file.

Terminating Threads in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_setschedparam Subroutine

### Purpose

Sets `schedpolicy` and `schedparam` attributes of a thread.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
#include <sys/sched.h>

int pthread_setschedparam (thread, schedpolicy, schedparam)
pthread_t thread;
int schedpolicy;
const struct sched_param *schedparam;
```

### Description

The `pthread_setschedparam` subroutine dynamically sets the `schedpolicy` and `schedparam` attributes of the thread `thread`. The `schedpolicy` attribute specifies the scheduling policy of the thread. The `schedparam` attribute specifies the scheduling parameters of a thread created with this attributes object. The `sched_priority` field of the `sched_param` structure contains the priority of the thread. It is an integer value.

If the target thread has system contention scope, the process must have root authority to set the scheduling policy to either `SCHED_FIFO` or `SCHED_RR`.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library. Otherwise, the `-D_THREAD_SAFE` compilation flag should be used, or the `cc_r` compiler used. In this case, the flag is automatically set.

This subroutine is part of the Base Operating System (BOS) Runtime. The implementation of this subroutine is dependent on the priority scheduling POSIX option. The priority scheduling POSIX option is implemented in the operating system.

## Parameters

<i>thread</i>	Specifies the target thread.
<i>schedpolicy</i>	Points to the <code>schedpolicy</code> attribute to set. It must have one of the following values: <b>SCHED_FIFO</b> Denotes first-in first-out scheduling. <b>SCHED_RR</b> Denotes round-robin scheduling. <b>SCHED_OTHER</b> Denotes the default operating system scheduling policy. It is the default value. If <i>schedpolicy</i> is <code>SCHED_OTHER</code> , then <i>sched_priority</i> must be in the range from 40 to 80, where 40 is the least favored priority and 80 is the most favored.
<i>schedparam</i>	<p><b>Note:</b> Priority of threads with a process contention scope and a <code>SCHED_OTHER</code> policy is controlled by the kernel; thus, setting the priority of such a thread has no effect. However, priority of threads with a system contention scope and a <code>SCHED_OTHER</code> policy can be modified. The modification directly affects the underlying kernel thread <code>nice</code> value.</p> <p>Points to where the scheduling parameters to set are stored. The <code>sched_priority</code> field must be in the range from 1 to 127, where 1 is the least favored priority, and 127 the most favored. If <i>schedpolicy</i> is <code>SCHED_OTHER</code>, then <i>sched_priority</i> must be in the range from 40 to 80, where 40 is the least favored priority and 80 is the most favored.</p> <p><b>Note:</b> Prior to AIX 5.3, users are not permitted to change the priority of a thread when setting its scheduling policy to <code>SCHED_OTHER</code>. In this case, the priority is managed directly by the kernel, and the only legal value that can be passed to <code>pthread_setschedparam</code> is <code>DEFAULT_PRIO</code>, which is defined in <code>pthread.h</code> as 1. All other passed values are ignored.</p> <p>Beginning with AIX 5.3, users can change the priority of a thread when setting its scheduling policy to <code>SCHED_OTHER</code>. The legal values that can be passed to <code>pthread_setschedparam</code> range from 40 to 80. Only privileged users can set a priority higher than 60. A value ranging from 1 to 39 provides the same priority as 40, and a value ranging from 81 to 127 provides the same priority as 80.</p>

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The `pthread_setschedparam` subroutine is unsuccessful if the following is true:

<b>EINVAL</b>	The <i>thread</i> or <i>schedparam</i> parameters are not valid.
<b>ENOSYS</b>	The priority scheduling POSIX option is not implemented.
<b>ENOTSUP</b>	The value of the <code>schedpolicy</code> or <code>schedparam</code> attributes are not supported.
<b>EPERM</b>	The target thread has insufficient permission to perform the operation or is already engaged in a mutex protocol.
<b>ESRCH</b>	The thread <i>thread</i> does not exist.

## Related Information

The **pthread\_getschedparam** (“pthread\_getschedparam Subroutine” on page 1255) subroutine, **pthread\_attr\_setschedparam** (“pthread\_attr\_setschedparam Subroutine” on page 1222) subroutine.

Threads Scheduling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## pthread\_setschedprio Subroutine

### Purpose

Dynamic thread scheduling parameters access (REALTIME THREADS).

### Syntax

```
#include <pthread.h>
```

```
int pthread_setschedprio(pthread_t thread, int prio);
```

### Description

The **pthread\_setschedprio()** function sets the scheduling priority for the thread whose thread ID is given by *thread* to the value given by *prio*. If a thread whose policy or priority has been modified by **pthread\_setschedprio()** is a running thread or is runnable, the effect on its position in the thread list depends on the direction of the modification as follows:

- If the priority is raised, the thread becomes the tail of the thread list.
- If the priority is unchanged, the thread does not change position in the thread list.
- If the priority is lowered, the thread becomes the head of the thread list.

Valid priorities are within the range returned by the **sched\_get\_priority\_max()** and **sched\_get\_priority\_min()**.

If the **pthread\_setschedprio()** function fails, the scheduling priority of the target thread remains unchanged.

### Rationale

The **pthread\_setschedprio()** function provides a way for an application to temporarily raise its priority and then lower it again, without having the undesired side-effect of yielding to other threads of the same priority. This is necessary if the application is to implement its own strategies for bounding priority inversion, such as priority inheritance or priority ceilings. This capability is especially important if the implementation does not support the **Thread Priority Protection** or **Thread Priority Inheritance** options; but even if those options are supported, this capability is needed if the application is to bound priority inheritance for other resources, such as semaphores.

The standard developers considered that, while it might be preferable conceptually to solve this problem by modifying the specification of **pthread\_setschedparam()**, it was too late to make such a change, because there might be implementations that would need to be changed. Therefore, this new function was introduced.

### Return Values

If successful, the **pthread\_setschedprio()** function returns 0; otherwise, an error number is returned to indicate the error.

## Error Codes

The `pthread_setschedprio()` function might fail if:

EINVAL	The value of <i>prio</i> is invalid for the scheduling policy of the specified thread.
ENOTSUP	An attempt was made to set the priority to an unsupported value.
EPERM	The caller does not have the appropriate permission to set the scheduling policy of the specified thread.
EPERM	The implementation does not allow the application to modify the priority to the value specified.
ESRCH	The value specified by <i>thread</i> does not refer to an existing thread.

The `pthread_setschedprio` function does not return an error code of [EINTR].

## Related Information

“`pthread_getschedparam` Subroutine” on page 1255, “`pthread_setschedparam` Subroutine” on page 1297.

The `pthread.h` file in *AIX 5L Version 5.3 Files Reference*.

---

## pthread\_sigmask Subroutine

### Purpose

Examines and changes blocked signals.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <signal.h>
```

```
int pthread_sigmask (how, set, oset)
int how;
const sigset_t *set;
sigset_t *oset;
```

### Description

Refer to `sigthreadmask` in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## pthread\_signal\_to\_cancel\_np Subroutine

### Purpose

Cancels the specified thread.

### Library

Threads Library (`libpthreads.a`)

### Syntax

```
#include <pthread.h>
```

```
int pthread_signal_to_cancel_np ( sigset, thread)
sigset_t *sigset;
pthread_t *thread;
```

## Description

The **pthread\_signal\_to\_cancel\_np** subroutine cancels the target thread *thread* by creating a handler thread. The handler thread calls the **sigwait** subroutine with the *sigset* parameter, and cancels the target thread when the **sigwait** subroutine returns. Successive calls to this subroutine override the previous ones.

### Notes:

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the *cc\_r* compiler used. In this case, the flag is automatically set.
2. The **pthread\_signal\_to\_cancel\_np** subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

## Parameters

*sigset* Specifies the set of signals to wait on.  
*thread* Specifies the thread to cancel.

## Return Values

Upon successful completion, 0 is returned. Otherwise, an error code is returned.

## Error Codes

The **pthread\_signal\_to\_cancel\_np** subroutine is unsuccessful if the following is true:

**EAGAIN** The handler thread cannot be created.  
**EINVAL** The *sigset* or *thread* parameters are not valid.

## Related Information

The **pthread\_cancel** (“pthread\_cancel Subroutine” on page 1231) subroutine, **sigwait** subroutine.

---

## pthread\_spin\_destroy or pthread\_spin\_init Subroutine

### Purpose

Destroys or initializes a spin lock object.

### Syntax

```
#include <pthread.h>

int pthread_spin_destroy(pthread_spinlock_t *lock);
int pthread_spin_init(pthread_spinlock_t *lock, int pshared);
```

### Description

The **pthread\_spin\_destroy** subroutine destroys the spin lock referenced by *lock* and releases any resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is

reinitialized by another call to the **pthread\_spin\_init** subroutine. The results are undefined if the **pthread\_spin\_destroy** subroutine is called when a thread holds the lock, or if this function is called with an uninitialized thread spin lock.

The **pthread\_spin\_init** subroutine allocates any resources required to use the spin lock referenced by lock and initializes the lock to an unlocked state.

If the Thread Process-Shared Synchronization option is supported and the value of *pshared* is PTHREAD\_PROCESS\_SHARED, the implementation shall permit the spin lock to be operated upon by any thread that has access to the memory where the spin lock is allocated, even if it is allocated in memory that is shared by multiple processes.

If the Thread Process-Shared Synchronization option is supported and the value of *pshared* is PTHREAD\_PROCESS\_PRIVATE, or if the option is not supported, the spin lock shall only be operated upon by threads created within the same process as the thread that initialized the spin lock. If threads of differing processes attempt to operate on such a spin lock, the behavior is undefined.

The results are undefined if the **pthread\_spin\_init** subroutine is called specifying an already initialized spin lock. The results are undefined if a spin lock is used without first being initialized.

If the **pthread\_spin\_init** subroutine function fails, the lock is not initialized and the contents of lock are undefined.

Only the object referenced by lock may be used for performing synchronization.

The result of referring to copies of that object in calls to the **pthread\_spin\_destroy** subroutine, **pthread\_spin\_lock** subroutine, **pthread\_spin\_trylock** subroutine, or the **pthread\_spin\_unlock** subroutine is undefined.

## Return Values

Upon successful completion, these functions shall return zero; otherwise, an error number shall be returned to indicate the error.

## Error Codes

<b>EBUSY</b>	The implementation has detected an attempt to initialize or destroy a spin lock while it is in use (for example, while being used in a <b>pthread_spin_lock</b> call) by another thread.
<b>EINVAL</b>	The value specified by the <i>lock</i> parameter is invalid.

The **pthread\_spin\_init** subroutine will fail if:

<b>EAGAIN</b>	The system lacks the necessary resources to initialize another spin lock.
<b>ENOMEM</b>	Insufficient memory exists to initialize the lock.

## Related Information

The “pthread\_spin\_lock or pthread\_spin\_trylock Subroutine,” “pthread\_spin\_unlock Subroutine” on page 1303.

---

## pthread\_spin\_lock or pthread\_spin\_trylock Subroutine

### Purpose

Locks a spin lock object.

## Syntax

```
#include <pthread.h>

int pthread_spin_lock(pthread_spinlock_t *lock);
int pthread_spin_trylock(pthread_spinlock_t *lock);
```

## Description

The **pthread\_spin\_lock** subroutine locks the spin lock referenced by the *lock* parameter. The calling thread shall acquire the lock if it is not held by another thread. Otherwise, the thread spins (that is, does not return from the **pthread\_spin\_lock** call) until the lock becomes available. The results are undefined if the calling thread holds the lock at the time the call is made. The **pthread\_spin\_trylock** subroutine locks the spin lock referenced by the *lock* parameter if it is not held by any thread. Otherwise, the function fails.

The results are undefined if any of these subroutines is called with an uninitialized spin lock.

## Return Values

Upon successful completion, these functions return zero; otherwise, an error number is returned to indicate the error.

## Error Codes

**EINVAL**                The value specified by the *lock* parameter does not refer to an initialized spin lock object.

The **pthread\_spin\_lock** subroutine fails if:

**EDEADLK**              The calling thread already holds the lock.

The **pthread\_spin\_trylock** subroutine fails if:

**EBUSY**                A thread currently holds the lock.

## Related Information

“pthread\_spin\_destroy or pthread\_spin\_init Subroutine” on page 1301, “pthread\_spin\_unlock Subroutine.”

---

## pthread\_spin\_unlock Subroutine

### Purpose

Unlocks a spin lock object.

### Syntax

```
#include <pthread.h>

int pthread_spin_unlock(pthread_spinlock_t *lock);
```

### Description

The **pthread\_spin\_unlock** subroutine releases the spin lock referenced by the *lock* parameter which was locked using the **pthread\_spin\_lock** subroutine or the **pthread\_spin\_trylock** subroutine. The results are undefined if the lock is not held by the calling thread. If there are threads spinning on the lock when the **pthread\_spin\_unlock** subroutine is called, the lock becomes available and an unspecified spinning thread shall acquire the lock.

The results are undefined if this subroutine is called with an uninitialized thread spin lock.

## Return Values

Upon successful completion, the **pthread\_spin\_unlock** subroutine returns zero; otherwise, an error number is returned to indicate the error.

## Error Codes

<b>EINVAL</b>	An invalid argument was specified.
<b>EPERM</b>	The calling thread does not hold the lock.

## Related Information

“pthread\_spin\_destroy or pthread\_spin\_init Subroutine” on page 1301, “pthread\_spin\_lock or pthread\_spin\_trylock Subroutine” on page 1302.

---

## pthread\_suspend\_np, pthread\_unsuspend\_np and pthread\_continue\_np Subroutine

### Purpose

Suspends and resume execution of the pthread specified by *thread*.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
```

```
pthread_t thread;  
int pthread_suspend_np(thread)  
int pthread_unsuspend_np (thread);  
int pthread_continue_np(thread);
```

### Description

The **pthread\_suspend\_np** subroutine immediately suspends the execution of the pthread specified by *thread*. On successful return from **pthread\_suspend\_np**, the suspended pthread is no longer executing. If **pthread\_suspend\_np** is called for a pthread that is already suspended, the pthread is unchanged and **pthread\_suspend\_np** returns successful.

Deadlock can occur if **pthread\_suspend\_np** is used with the following pthread functions.

```
pthread_getrusage_np  
pthread_cancel  
pthread_detach  
pthread_join  
pthread_getunique_np  
pthread_join_np  
pthread_setschedparam  
pthread_getschedparam  
pthread_kill
```

To prevent deadlock, **PTHREAD\_SUSPENDIBLE=ON** should be set.



The **pthread\_unsuspend\_np** routine decrements the suspend count and once the count is zero, the routine resumes the execution of a suspended pthread. If **pthread\_unsuspend\_np** is called for a pthread that is not suspended, the pthread is unchanged and **pthread\_unsuspend\_np** returns successful.

The **pthread\_continue\_np** routine clears the suspend count and resumes the execution of a suspended pthread. If **pthread\_continue\_np** is called for a pthread that is not suspended, the pthread is unchanged and **pthread\_continue\_np** returns successful.

A suspended pthread will not be awakened by a signal. The signal stays pending until the execution of pthread is resumed by **pthread\_continue\_np**.

**Note:** Using **pthread\_suspend\_np** should only be used by advanced users because improper use of this subcommand can lead to application deadlock or the target thread may be suspended holding application locks.

## Parameters

*thread*                Specifies the target thread.

## Return Values

Zero is returned when successful. A nonzero value indicates an error.

## Error Codes

If any of the following conditions occur, **pthread\_suspend\_np**, **pthread\_unsuspend\_np** and **pthread\_continue\_np** fail and return the corresponding value:

**ESRCH**                The target thread specified by *thread* attribute cannot be found in the current process.

---

## pthread\_unlock\_global\_np Subroutine

### Purpose

Unlocks the global mutex.

### Library

Threads Library (**libpthreads.a**)

### Syntax

```
#include <pthread.h>
void pthread_unlock_global_np ()
```

### Description

The **pthread\_unlock\_global\_np** subroutine unlocks the global mutex when each call to the **pthread\_lock\_global\_np** subroutine is matched by a call to this routine. For example, if a thread called the **pthread\_lock\_global\_np** three times, the global mutex is unlocked after the third call to the **pthread\_unlock\_global\_np** subroutine.

If no threads are waiting for the global mutex, it becomes unlocked with no current owner. If one or more threads are waiting to lock the global mutex, exactly one thread returns from its call to the **pthread\_lock\_global\_np** subroutine.

**Notes:**

1. The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the cc\_r compiler used. In this case, the flag is automatically set.
2. The **pthread\_unlock\_global\_np** subroutine is not portable.

This subroutine is not POSIX compliant and is provided only for compatibility with DCE threads. It should not be used when writing new applications.

**Related Information**

The **pthread\_lock\_global\_np** (“pthread\_lock\_global\_np Subroutine” on page 1266) subroutine.

Using Mutexes in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

**pthread\_yield Subroutine****Purpose**

Forces the calling thread to relinquish use of its processor.

**Library**

Threads Library (**libpthreads.a**)

**Syntax**

```
#include <pthread.h>
void pthread_yield ()
```

**Description**

The **pthread\_yield** subroutine forces the calling thread to relinquish use of its processor, and to wait in the run queue before it is scheduled again. If the run queue is empty when the **pthread\_yield** subroutine is called, the calling thread is immediately rescheduled.

If the thread has global contention scope (**PTHREAD\_SCOPE\_SYSTEM**), calling this subroutine acts like calling the **yield** subroutine. Otherwise, another local contention scope thread is scheduled.

The **pthread.h** header file must be the first included file of each source file using the threads library. Otherwise, the **-D\_THREAD\_SAFE** compilation flag should be used, or the cc\_r compiler used. In this case, the flag is automatically set.

**Related Information**

The **yield** subroutine and the **sched\_yield** subroutine.

Threads Scheduling in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

Threads Library Options in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## ptrace, ptracex, ptrace64 Subroutine

### Purpose

Traces the execution of another process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/reg.h>
#include <sys/ptrace.h>
#include <sys/ldr.h>
int ptrace ( Request, Identifier, Address, Data, Buffer)
int Request;
int Identifier;
int *Address;
int Data;
int *Buffer;
```

```
int ptracex ( request, identifier, addr, data, buff)
int request;
int identifier;
long long addr;
int data;
int *buff;
```

```
int ptrace64 ( request, identifier, addr, data, buff)
int request;
long long identifier;
long long addr;
int data;
int *buff;
```

### Description

The **ptrace** subroutine allows a 32-bit process to trace the execution of another process. The **ptrace** subroutine is used to implement breakpoint debugging.

A debugged process runs normally until it encounters a signal. Then it enters a stopped state and its debugging process is notified with the **wait** subroutine.

**Exception:** If the process encounters the **SIGTRAP** signal, a signal handler for **SIGTRAP** exists, and fast traps (“Fast Trap Instructions” on page 1308) have been enabled for the process, then the signal handler is called and the debugger is not notified. This exception only applies to AIX 4.3.3 and later releases.

While the process is in the stopped state, the debugger examines and modifies the memory image of the process being debugged by using the **ptrace** subroutine. For multi-threaded processes, the **getthrds** (“getthrds Subroutine” on page 443) subroutine identifies each kernel thread in the debugged process. Also, the debugging process can cause the debugged process to terminate or continue, with the possibility of ignoring the signal that caused it to stop.

As a security measure, the **ptrace** subroutine inhibits the set-user-ID facility on subsequent **exec** subroutines.

(This paragraph only applies to AIX 4.3.2 and later releases.) When a process running under **ptrace** control calls **load** or **unload**, the debugger is notified and the **W\_SLWTEd** flag is set in the status returned by **wait**. (A 32-bit process calling **loadbind** is stopped as well.) If the process being debugged has added modules in the shared library to its address space, the modules are added to the process's private copy of the shared library segments. If shared library modules are removed from a process's address space, the modules are deleted from the process's private copy of the library text segment by freeing the pages that contain the module. No other changes to the segment are made, and existing breakpoints do not have to be reinserted.

To allow a debugger to generate code more easily (in order to handle fast trap instructions, for example), memory from the end of the main program up to the next segment boundary can be modified. That memory is read-only to the process but can be modified by the debugger.

When a process being traced forks, the child process is initialized with the unmodified main program and shared library segment, effectively removing breakpoints in these segments in the child process. If multiprocess debugging is enabled, new copies of the main program and shared library segments are made. Modifications to privately loaded modules, however, are not affected by a fork. These breakpoints will remain in the child process, and if these breakpoints are run, a **SIGTRAP** signal is generated and delivered to the process.

If a traced process initiates an **exec** subroutine, the process stops before executing the first instruction of the new image and returns the **SIGTRAP** signal.

**Note:** The **ptrace** and **ptracex** subroutines are not supported in 64-bit mode.

### Fast Trap Instructions

Sometimes, allowing the process being debugged to handle certain trap instructions is useful, instead of causing the process to stop and notify the debugger. You can use this capability to patch running programs or programs whose source codes are not available. For a process to use this capability, you must enable fast traps, which requires you to make a **ptrace** call from a debugger on behalf of the process.

To let a process handle fast traps, a debugger uses the **ptrace (PT\_SET, pid, 0, PTFLAG\_FAST\_TRAP, 0)** subroutine call. Cancel this capability with the **ptrace (PT\_CLEAR, pid, 0, PTFLAG\_FAST\_TRAP, 0)** subroutine call. If a process is able to handle fast traps when the debugger detaches, the fast trap capability remains in effect. Consequently, when another debugger attaches to that process, fast trap processing is still enabled. When no debugger is attached to a process, **SIGTRAP** signals are handled in the same manner, regardless of whether fast traps are enabled.

A fast trap instruction is an unconditional *trap immediate* instruction in the form `twi 14,r13,0xNXXX`. This instruction has the binary form `0x0ddfNXXX`, where N is a hex digit  $\geq 8$  and XXX are any three hex digits. By using different values of `0xNXXX`, a debugger can generate different fast trap instructions, allowing a signal handler to quickly determine how to handle the signal. (The fast trap instruction is defined by the macro **\_PTRACE\_FASTTRAP**. The **\_PTRACE\_FASTTRAP\_MASK** macro can be used to check whether a trap is a fast trap.)

Usually, a fast trap instruction is treated like any other trap instruction. However, if a process has a signal handler for **SIGTRAP**, the signal is not blocked, and the fast trap capability is enabled, then the signal handler is called and the debugger is not notified.

A signal handler can logically AND the trap instruction with **\_PTRACE\_FASTTRAP\_NUM (0x7FFF)** to obtain an integer identifying which trap instruction was run.

### Fast data watchpoint

In AIX 5.3 ML5 and later, **ptrace** supports the ability to enable fast watchpoint trap handling. This is similar to fast trap instruction handling in that when it is enabled. Processes that have a signal handler for

**SIGTRAP** will have the handler called when a watchpoint trap is encountered. In the **SIGTRAP** signal handler, the traced process can detect a fast watchpoint trap by checking the **SI\_FAST\_WATCH** in the **\_si\_flags** of the **siginfo\_t** that is passed to the handler. The fast watchpoint handling employs trap-after semantics, which means that the store to the watched location is completed before calling the trap handler, so the instruction address pointer in the signal context that is passed to the handler will point to the instruction following the instruction that caused the trap.

## Thread-level tracing

In AIX 5.3 ML5 and later, **ptrace** supports setting breakpoints and watchpoints per-thread for system scope (1:1) threads. With these, the tracing process (debugger) is only notified when the specific thread of interest has encountered a trap. This provides an efficient means for debuggers to trace individual threads of interest since it doesn't have to filter "false hit" notifications. See the **PTT\_WATCH**, **PTT\_SET\_TRAP**, and **PTT\_CLEAR\_TRAP** request types below for the usage description.

The **ptrace** programming model remains unchanged with thread-level breakpoints and watchpoints in that the attachment is still done at the process level, and the target process stops and notifies the tracing process upon encountering a trap. The tracing process can detect that the traced process has stopped for a thread-level trap by checking the **TTHRDRTRAP** flag (in **ti\_flag2**) of the stopping thread (the thread with **TTRCSIG** set in **ti\_flag**). These flags can be checked by calling **getthrds64** on the target process.

Other behaviors that are specific to thread-level tracing:

Thread-level breakpoints

- Clear automatically when all threads for which the breakpoint is active have terminated.
- Not supported for multiprocess debugging (**PT\_MULTI**). They are cleared upon **fork** and **exec**.

Thread-level watchpoints

- Newly created threads inherit the process-level watch location.
- Not inherited across **fork** and **exec**.

## For the 64-bit Process

Use **ptracex** where the debuggee is a 64-bit process and the operation requested uses the third (*Address*) parameter to reference the debuggee's address space or is sensitive to register size. Note that **ptracex** and **ptrace64** will also support 32-bit debuggees.

If returning or passing an **int** doesn't work for a 64-bit debuggee (for example, **PT\_READ\_GPR**), the buffer parameter takes the address for the result. Thus, with the **ptracex** subroutine, **PT\_READ\_GPR** and **PT\_WRITE\_GPR** take a pointer to an 8 byte area representing the register value.

In general, **ptracex** supports all the calls that **ptrace** does when they are modified for any that are extended for 64-bit addresses (for example, GPRs, LR, CTR, IAR, and MSR). Anything whose size increases for 64-bit processes must be allowed for in the obvious way (for example, **PT\_REGSET** must be an array of long longs for a 64-bit debuggee).

## Parameters

*Request*

Determines the action to be taken by the **ptrace** subroutine and has one of the following values:

### **PT\_ATTACH**

This request allows a debugging process to attach a current process and place it into trace mode for debugging. This request cannot be used if the target process is already being traced. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Address*, *Data*, and *Buffer* parameters are ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one of the following codes:

**ESRCH**

*Process* ID is not valid; the traced process is a kernel process; the process is currently being traced; or, the debugger or traced process already exists.

**EPERM**

Real or effective user ID of the debugger does not match that of the traced process, or the debugger does not have root authority.

**EINVAL**

The debugger and the traced process are the same.

**PT\_CLEAR**

This request clears an internal flag or capability. The *Data* parameter specifies which flags to clear. The following flag can be cleared:

**PTFLAG\_FAST\_TRAP**

Disables the special handling of a fast trap instruction (“Fast Trap Instructions” on page 1308). This allows all fast trap instructions causing an interrupt to generate a **SIGTRAP** signal.

The *Identifier* parameter specifies the process ID of the traced process. The *Address* parameter, *Buffer* parameter, and the unused bits in the *Data* parameter are reserved for future use and should be set to 0.

**PTFLAG\_FAST\_WATCH**

Enables fast watchpoint trap handling. When a watchpoint trap occurs in a process that has a signal handler for **SIGTRAP**, and the process has fast watchpoints enabled, the signal handler will be called instead of notifying the tracing process.

**PTT\_CLEAR\_TRAP**

This request type clears thread-level breakpoints.

The *Identifier* parameter is a valid kernel thread ID in the target process (-1 for all). The *Address* parameter is the address of the breakpoint. The *Data* parameter must be 0. The *Buffer* parameter must be NULL.

If the request is unsuccessful, -1 is returned and the **errno** global variable is set to one of the following:

**ESRCH**

The *Identifier* parameter does not refer to a valid kernel thread in the target process, or no breakpoint was found for the target thread at the given *Address*.

**EINVAL**

The *Data* parameter was non-zero or *Buffer* was non-NULL.

**PT\_CONTINUE**

This request allows the process to resume execution. If the *Data* parameter is 0, all pending signals, including the one that caused the process to stop, are concealed before the process resumes execution. If the *Data* parameter is a valid signal number, the process resumes execution as if it had received that signal. If the *Address* parameter equals 1, the execution continues from where it stopped. If the *Address* parameter is not 1, it is assumed to be the address at which the process should resume execution. Upon

successful completion, the value of the *Data* parameter is returned to the debugging process. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The signal to be sent to the traced process is not a valid signal number.

**Note:** For the **PT\_CONTINUE** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the resume address needs 64 bits.

## PTT\_CONTINUE

This request asks the scheduler to resume execution of the kernel thread specified by *Identifier*. This kernel thread must be the one that caused the exception. The *Data* parameter specifies how to handle signals:

- If the *Data* parameter is 0, the kernel thread which caused the exception will be resumed as if the signal never occurred.
- If the *Data* parameter is a valid signal number, the kernel thread which caused the exception will be resumed as if it had received that signal.

The *Address* parameter specifies where to resume execution:

- If the *Address* parameter is 1, execution resumes from the address where it stopped.
- If the *Address* parameter contains an address value other than 1, execution resumes from that address.

The *Buffer* parameter should point to a PTTTHREADS structure, which contains a list of kernel thread identifiers to be started. This list should be NULL terminated if it is smaller than the maximum allowed.

On successful completion, the value of the *Data* parameter is returned to the debugging process. On unsuccessful completion, the value -1 is returned, and the **errno** global variable is set as follows:

### EINVAL

The *Identifier* parameter names the wrong kernel thread.

**EIO** The signal to be sent to the traced kernel thread is not a valid signal number.

### ESRCH

The *Buffer* parameter names an invalid kernel thread. Each kernel thread in the list must be stopped and belong to the same process as the kernel thread named by the *Identifier* parameter.

**Note:** For the **PTT\_CONTINUE** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the resume address needs 64 bits.

## PT\_DETACH

This request allows a debugged process, specified by the *Identifier* parameter, to exit trace mode. The process then continues running, as if it had received the signal whose number is contained in the *Data* parameter. The process is no longer traced and does not process any further **ptrace** calls. The *Address* and *Buffer* parameters are ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** Signal to be sent to the traced process is not a valid signal number.

## PT\_GET\_UKEY

This request reads the user-key assigned to a specific effective address indicated by the *address* parameter into the location pointed to the *buffer* parameter. The process ID of the traced process must be passed in the *identifier* parameter. The *data* parameter is ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

### ENOSYS

Process is not user-key aware.

## PT\_KILL

This request allows the process to terminate the same way it would with an **exit** subroutine.

## PT\_LDINFO

This request retrieves a description of the object modules that were loaded by the debugged process. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored. The *Address* parameter specifies the location where the loader information is copied. The *Data* parameter specifies the size of this area. The loader information is retrieved as a linked list of **ld\_info** structures. The first element of the list corresponds to the main executable module. The **ld\_info** structures are defined in the **/usr/include/sys/ldr.h** file. The linked list is implemented so that the `ldinfo_next` field of each element gives the offset of the next element from this element. The `ldinfo_next` field of the last element has the value 0.

Each object module reported is opened on behalf of the debugger process. The file descriptor for an object module is saved in the `ldinfo_fd` field of the corresponding **ld\_info** structure. The debugger process is responsible for managing the files opened by the **ptrace** subroutine.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

### ENOMEM

Either the area is not large enough to accommodate the loader information, or there is not enough memory to allocate an equivalent buffer in the kernel.

**Note:** For the **PT\_LDINFO** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the source address needs 64 bits.

## PT\_LDXINFO

This request is similar to the **PT\_LDINFO** request. A linked list of **ld\_xinfo** structures is returned instead of a list of **ld\_info** structures. The first element of the list corresponds to the main executable module. The **ld\_xinfo** structures are defined in the **/usr/include/sys/ldr.h** file. The linked list is implemented so that the `ldinfo_next` field of each element gives the offset of the next element from this element. The `ldinfo_next` field of the last element has the value 0.

Each object module reported is opened on behalf of the debugger process. The file descriptor for an object module is saved in the `ldinfo_fd` field of the corresponding **ld\_xinfo** structure. The debugger process is responsible for managing the files opened by the **ptrace** subroutine.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:



## ENOMEM

Either the area is not large enough to accommodate the loader information, or there is not enough memory to allocate an equivalent buffer in the kernel.

**Note:** For the **PT\_LDXINFO** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the source address needs 64 bits.

## PT\_MULTI

This request turns multiprocess debugging mode on and off, to allow debugging to continue across **fork** and **exec** subroutines. A 0 value for the *Data* parameter turns multiprocess debugging mode off, while all other values turn it on. When multiprocess debugging mode is in effect, any **fork** subroutine allows both the traced process and its newly created process to trap on the next instruction. If a traced process initiated an **exec** subroutine, the process stops before executing the first instruction of the new image and returns the **SIGTRAP** signal. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Address* and *Buffer* parameters are ignored.

Also, when multiprocess debugging mode is enabled, the following values are returned from the **wait** subroutine:

### W\_SEWTED

Process stopped during execution of the **exec** subroutine.

### W\_SFWTED

Process stopped during execution of the **fork** subroutine.

## PT\_READ\_BLOCK

This request reads a block of data from the debugged process address space. The *Address* parameter points to the block of data in the process address space, and the *Data* parameter gives its length in bytes. The value of the *Data* parameter must not be greater than 1024. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter points to the location in the debugging process address space where the data is copied. Upon successful completion, the **ptrace** subroutine returns the value of the *Data* parameter.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one of the following codes:

**EIO** The *Data* parameter is less than 1 or greater than 1024.

**EIO** The *Address* parameter is not a valid pointer into the debugged process address space.

### EFAULT

The *Buffer* parameter does not point to a writable location in the debugging process address space.

**Note:** For the **PT\_READ\_BLOCK** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the source address needs 64 bits.

## PT\_READ\_FPR

This request stores the value of a floating-point register into the location pointed to by the *Address* parameter. The *Data* parameter specifies the floating-point register, defined in the **sys/reg.h** file for the machine type on which the process is run. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The *Data* parameter is not a valid floating-point register. The *Data* parameter must be in the range 256-287.

#### **PTT\_READ\_FPRS**

This request writes the contents of the 32 floating point registers to the area specified by the *Address* parameter. This area must be at least 256 bytes long. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

#### **PTT\_READ\_FPSCR\_HI**

This request writes the contents of the upper 32-bits of the FPSCR register to the area specified by the *Address* parameter. This area must be at least 4 bytes long. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

#### **PTT\_WRITE\_FPSCR\_HI**

This request updates the contents of the upper 32-bits of the FPSCR register with the value specified in the area designated by the *Address* parameter. This area must be at least 4 bytes long. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

#### **PT\_READ\_GPR**

This request returns the contents of one of the general-purpose or special-purpose registers of the debugged process. The *Address* parameter specifies the register whose value is returned. The value of the *Address* parameter is defined in the **sys/reg.h** file for the machine type on which the process is run. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Data* and *Buffer* parameters are ignored. The buffer points to long long target area.

**Note:** If **ptracex** or **ptrace64** with a 64-bit debuggee is used for this request, the register value is instead returned to the 8-byte area pointed to by the buffer pointer.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The *Address* is not a valid general-purpose or special-purpose register. The *Address* parameter must be in the range 0-31 or 128-136.

#### **PTT\_READ\_GPRS**

This request writes the contents of the 32 general purpose registers to the area specified by the *Address* parameter. This area must be at least 128 bytes long.

**Note:** If **ptracex** or **ptrace64** are used with a 64-bit debuggee for the **PTT\_READ\_GPRS** request, there must be at least a 256 byte target area. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

#### **PT\_READ\_I or PT\_READ\_D**

These requests return the word-aligned address in the debugged process address space specified by the *Address* parameter. On all machines currently supported by AIX Version 4, the **PT\_READ\_I** and **PT\_READ\_D** instruction and data requests can be used with equal results. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Data* parameter is ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The *Address* is not word-aligned, or the *Address* is not valid. User blocks, kernel segments, and kernel extension segments are not considered as valid addresses.

**Note:** For the **PT\_READ\_I** or the **PT\_READ\_D** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the source address needs 64 bits.

### **PTT\_READ\_SPRS**

This request writes the contents of the special purpose registers to the area specified by the *Address* parameter, which points to a **ptsprs** structure. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

**Note:** For the **PTT\_READ\_SPRS** request, use **ptracex** or **ptrace64** with the 64-bit debuggee because the new **ptxsprs** structure must be used.

### **PTT\_READ\_UKEYSET**

This request reads the active user-key-set for the specified thread whose thread ID is specified by the *identifier* parameter into the location pointed to the *buffer* parameter. The *address* and *data* parameters are ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

#### **ENOSYS**

Process is not user-key aware.

### **PTT\_READ\_VEC**

This request reads the vector register state of the specified thread. The data format is a **\_\_vmx\_context\_t** structure that contains the 32 vector registers, in addition to the VSCR and VRSAVE registers.

### **PT\_REATT**

This request allows a new debugger, with the proper permissions, to trace a process that was already traced by another debugger. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Address*, *Data*, and *Buffer* parameters are ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one the following codes:

#### **ESRCH**

The *Identifier* is not valid; or the traced process is a kernel process.

#### **EPERM**

Real or effective user ID of the debugger does not match that of the traced process, or the debugger does not have root authority.

#### **EINVAL**

The debugger and the traced process are the same.

### **PT\_REGSET**

This request writes the contents of all 32 general purpose registers to the area specified by the *Address* parameter. This area must be at least 128 bytes for the 32-bit debuggee or 256 bytes for the 64-bit debuggee. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Data* and *Buffer* parameters are ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

#### **EIO**

The *Address* parameter points to a location outside of the allocated address space of the process.

**Note:** For the **PT\_REGSET** request, use **ptracex** or **ptrace64** with the 64-bit debuggee because 64-bit registers requiring 256 bytes are returned.

## PT\_SET

This request sets an internal flag or capability. The *Data* parameter indicates which flags are set. The following flag can be set:

### PTFLAG\_FAST\_TRAP

Enables the special handling of a fast trap instruction (“Fast Trap Instructions” on page 1308). When a fast trap instruction is run in a process that has a signal handler for **SIGTRAP**, the signal handler will be called even if the process is being traced.

The *Identifier* parameter specifies the process ID of the traced process. The *Address* parameter, *Buffer* parameter, and the unused bits in the *Data* parameter are reserved for future use and should be set to 0.

## PTT\_SET\_TRAP

This request type sets thread-level breakpoints.

The *Identifier* parameter is a valid kernel ID in the target process. The *Address* parameter is the address in the target process for the breakpoint. The *Data* parameter is the length of data in *Buffer*, it must be 4. The *Buffer* parameter is a pointer to trap instruction to be written.

The system call will not evaluate the contents of the buffer for this request, but by convention, it should contain a single trap instruction.

If the request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one of the following:

### ENOMEM

Could not allocate kernel memory.

### ESRCH

The *Identifier* parameter does not refer to a valid kernel thread in the target process.

### EIO

The *Address* parameter does not point to a writable location in the address space of the target process.

### EINVAL

*Data* parameter was not 4, or the target thread already has a breakpoint set at *Address*.

### EFAULT

The *Buffer* parameter does not point to a readable location in the caller’s address space.

## PT\_TRACE\_ME

This request must be issued by the debugged process to be traced. Upon receipt of a signal, this request sets the process trace flag, placing the process in a stopped state, rather than the action specified by the **sigaction** subroutine. The *Identifier*, *Address*, *Data*, and *Buffer* parameters are ignored. Do not issue this request if the parent process does not expect to trace the debugged process.

As a security measure, the **ptrace** subroutine inhibits the set-user-ID facility on subsequent **exec** subroutines, as shown in the following example:

```
if((childpid = fork()) == 0)
{ /* child process */
  ptrace(PT_TRACE_ME,0,0,0,0);
```

```

        execvp(          )/* your favorite exec*/
    }
else
{
    /* parent
    /* wait for child to stop
        rc = wait(status)
    */
}

```

**Note:** This is the only request that should be performed by the child. The parent should perform all other requests when the child is in a stopped state.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

#### **ESRCH**

Process is debugged by a process that is not its parent.

### **PT\_WATCH**

This request allows to have a watchpoint on the memory region specified when the debugged process changes the content at the specified memory region.

The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored. The *Address* parameter specifies beginning of the memory region to be watched. To clear the watchpoint the *Address* parameter must be NULL. The *Data* parameter specifies the size of the memory region.

Watchpoints are supported only on the hardware POWER630, POWER5 and POWER6. Currently the size of the memory region, that is, the parameter *Data* must be 8 because only 8 byte watchpoint is supported at the hardware level.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

#### **EPERM**

If the hardware does not support watchpoints or if specified *Identifier* is not valid Process ID.

**EIO** If the specified *Address* is not double word aligned.

#### **EINVAL**

If the specified *Data* is not 8.

### **PTT\_WATCH**

This request sets and clears thread-level watchpoints.

The *Identifier* parameter is a valid kernel thread ID in the target process ( -1 for all). The *Address* parameter is the double-worded aligned address to watch. A value of 0 clears the watchpoint. The *Data* parameter must be 0 (clear) or 8 (set). The *Buffer* parameter must be NULL.

If the request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one of the following:

#### **ESRCH**

The *Identifier* parameter does not refer to a valid kernel thread in the target process.

#### **EPERM**

The hardware watchpoint facility is not supported on the platform.

**EIO** The requested *Address* is not a valid, double-worded aligned address in target process address space, or the *Address* is non-zero and *Data* is not 8

## PT\_WRITE\_BLOCK

This request writes a block of data into the debugged process address space. The *Address* parameter points to the location in the process address space to be written into. The *Data* parameter gives the length of the block in bytes, and must not be greater than 1024. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter points to the location in the debugging process address space where the data is copied. Upon successful completion, the value of the *Data* parameter is returned to the debugging process.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to one of the following codes:

**EIO** The *Data* parameter is less than 1 or greater than 1024.

**EIO** The *Address* parameter is not a valid pointer into the debugged process address space.

### EFAULT

The *Buffer* parameter does not point to a readable location in the debugging process address space.

**Note:** For the **PT\_WRITE\_BLOCK** request, use **ptracex** or **ptrace64** with the 64-bit debuggee because 64-bit registers requiring 256 bytes are returned.

## PT\_WRITE\_FPR

This request sets the floating-point register specified by the *Data* parameter to the value specified by the *Address* parameter. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The *Data* parameter is not a valid floating-point register. The *Data* parameter must be in the range 256-287.

## PTT\_WRITE\_FPRS

This request updates the contents of the 32 floating point registers with the values specified in the area designated by the *Address* parameter. This area must be at least 256 bytes long. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

## PT\_WRITE\_GPR

This request stores the value of the *Data* parameter in one of the process general-purpose or special-purpose registers. The *Address* parameter specifies the register to be modified. Upon successful completion, the value of the *Data* parameter is returned to the debugging process. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored.

**Note:** If **ptracex** or **ptrace64** are used with a 64-bit debuggee for the **PT\_WRITE\_GPR** request, the new register value is NOT passed via the *Data* parameter, but is instead passed via the 8-byte area pointed to by the buffer parameter.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The *Address* parameter is not a valid general-purpose or special-purpose register. The *Address* parameter must be in the range 0-31 or 128-136.

### PTT\_WRITE\_GPRS

This request updates the contents of the 32 general purpose registers with the values specified in the area designated by the *Address* parameter. This area must be at least 128 bytes long. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

**Note:** For the **PTT\_WRITE\_GPRS** request, use **ptracex** or **ptrace64** with the 64-bit debuggee because 64-bit registers requiring 256 bytes are returned. The buffer points to long long source area.

### PT\_WRITE\_I or PT\_WRITE\_D

These requests write the value of the *Data* parameter into the address space of the debugged process at the word-aligned address specified by the *Address* parameter. On all machines currently supported by AIX Version 4, instruction and data address spaces are not separated. The **PT\_WRITE\_I** and **PT\_WRITE\_D** instruction and data requests can be used with equal results. Upon successful completion, the value written into the address space of the debugged process is returned to the debugging process. The *Identifier* parameter is interpreted as the process ID of the traced process. The *Buffer* parameter is ignored.

If this request is unsuccessful, a value of -1 is returned and the **errno** global variable is set to the following code:

**EIO** The *Address* parameter points to a location in a pure procedure space and a copy cannot be made; the *Address* is not word-aligned; or, the *Address* is not valid. User blocks, kernel segments, and kernel extension segments are not considered valid addresses.

**Note:** For the or **PT\_WRITE\_I** or **PT\_WRITE\_D** request, use **ptracex** or **ptrace64** with a 64-bit debuggee because the target address needs 64 bits.

### PTT\_WRITE\_SPRS

This request updates the special purpose registers with the values in the area specified by the *Address* parameter, which points to a **ptsprs** structure. The *Identifier* parameter specifies the traced kernel thread. The *Data* and *Buffer* parameters are ignored.

#### *Identifier*

Determined by the value of the *Request* parameter.

#### *Address*

Determined by the value of the *Request* parameter.

*Data* Determined by the value of the *Request* parameter.

*Buffer* Determined by the value of the *Request* parameter.

**Note:** For the **PTT\_READ\_SPRS** request, use **ptracex** or **ptrace64** with the 64-bit debuggee because the new **ptsprs** structure must be used.

### PTT\_WRITE\_VEC

This request writes the vector register state of the specified thread. The data format is a **\_\_vmx\_context\_t** structure that contains the 32 vector registers, in addition to the VSCR and VRSAVE registers.

## Error Codes

The `ptrace` subroutine is unsuccessful when one of the following is true:

<b>EFAULT</b>	The <i>Buffer</i> parameter points to a location outside the debugging process address space.
<b>EINVAL</b>	The debugger and the traced process are the same; or the <i>Identifier</i> parameter does not identify the thread that caused the exception.
<b>EIO</b>	The <i>Request</i> parameter is not one of the values listed, or the <i>Request</i> parameter is not valid for the machine type on which the process is run.
<b>ENOMEM</b>	Either the area is not large enough to accommodate the loader information, or there is not enough memory to allocate an equivalent buffer in the kernel.
<b>ENXIO</b>	The target thread has not referenced the VMX unit and is not currently a VMX thread.
<b>EPERM</b>	The <i>Identifier</i> parameter corresponds to a kernel thread which is stopped in kernel mode and whose computational state cannot be read or written.
<b>ESRCH</b>	The <i>Identifier</i> parameter identifies a process or thread that does not exist, that has not run a <code>ptrace</code> call with the <code>PT_TRACE_ME</code> request, or that is not stopped.

For `ptrace`: If the debuggee is a 64-bit process, the options that refer to GPRs or SPRs fail with `errno = EIO`, and the options that specify addresses are limited to 32-bits.

For `ptracex` or `ptrace64`: If the debuggee is a 32-bit process, the options that refer to GPRs or SPRs fail with `errno = EIO`, and the options that specify addresses in the debuggee's address space that are larger than  $2^{32} - 1$  fail with `errno` set to `EIO`.

Also, the options `PT_READ_U` and `PT_WRITE_U` are not supported if the debuggee is a 64-bit program (`errno = ENOTSUP`).

## Related Information

The “`exec: execl, execl, execlp, execv, execve, execvp, or exact Subroutine`” on page 236, “`getprocs Subroutine`” on page 415, “`getthrds Subroutine`” on page 443, and “`load and loadAndInit Subroutines`” on page 727.

The `sigaction` subroutine, `unload` subroutine, and `wait`, `waitpid`, or `wait3` subroutine in *AIX 5L Version 5.3 Technical Reference: Base Operating System and Extensions Volume 2*.

The `dbx` command in *AIX 5L Version 5.3 Commands Reference, Volume 2*.

The `sys/ldr.h` file.

---

## ptsname Subroutine

### Purpose

Returns the name of a pseudo-terminal device.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
```

```
char *ptsname ( FileDescriptor )  
int FileDescriptor
```



## Description

The **ptsname** subroutine gets the path name of the slave pseudo-terminal associated with the master pseudo-terminal device defined by the *FileDescriptor* parameter.

## Parameters

*FileDescriptor* Specifies the file descriptor of the master pseudo-terminal device

## Return Values

The **ptsname** subroutine returns a pointer to a string containing the null-terminated path name of the pseudo-terminal device associated with the file descriptor specified by the *FileDescriptor* parameter. A null pointer is returned and the **errno** global variable is set to indicate the error if the file descriptor does not describe a pseudo-terminal device in the */dev* directory.

## Files

*/dev/\** Terminal device special files.

## Related Information

The **ttyname** subroutine.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## putc, putchar, fputc, or putw Subroutine

### Purpose

Writes a character or a word to a stream.

### Library

Standard I/O Package (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
int putc ( Character, Stream)
int Character;
FILE *Stream;
```

```
int putchar (Character)
int Character;
```

```
int fputc (Character, Stream)
int Character;
FILE *Stream;
```

```
int putw ( Word, Stream)
int Word;
FILE *Stream;
```

## Description

The **putc** and **putchar** macros write a character or word to a stream. The **fputc** and **putw** subroutines serve similar purposes but are true subroutines.

The **putc** macro writes the character *Character* (converted to an **unsigned char** data type) to the output specified by the *Stream* parameter. The character is written at the position at which the file pointer is currently pointing, if defined.

The **putchar** macro is the same as the **putc** macro except that **putchar** writes to the standard output.

The **fputc** subroutine works the same as the **putc** macro, but **fputc** is a true subroutine rather than a macro. It runs more slowly than **putc**, but takes less space per invocation.

Because **putc** is implemented as a macro, it incorrectly treats a *Stream* parameter with side effects, such as **putc(C, \*f++)**. For such cases, use the **fputc** subroutine instead. Also, use **fputc** whenever you need to pass a pointer to this subroutine as a parameter to another subroutine.

The **putc** and **putchar** macros have also been implemented as subroutines for ANSI compatibility. To access the subroutines instead of the macros, insert **#undef putc** or **#undef putchar** at the beginning of the source file.

The **putw** subroutine writes the word (**int** data type) specified by the *Word* parameter to the output specified by the *Stream* parameter. The word is written at the position at which the file pointer, if defined, is pointing. The size of a word is the size of an integer and varies from machine to machine. The **putw** subroutine does not assume or cause special alignment of the data in the file.

After the **fputcw**, **putwc**, **fputc**, **putc**, **fputs**, **puts**, or **putw** subroutine runs successfully, and before the next successful completion of a call either to the **fflush** or **fclose** subroutine on the same stream or to the **exit** or **abort** subroutine, the `st_ctime` and `st_mtime` fields of the file are marked for update.

Because of possible differences in word length and byte ordering, files written using the **putw** subroutine are machine-dependent, and may not be readable using the **getw** subroutine on a different processor.

With the exception of **stderr**, output streams are, by default, buffered if they refer to files, or line-buffered if they refer to terminals. The standard error output stream, **stderr**, is unbuffered by default, but using the **freopen** subroutine causes it to become buffered or line-buffered. Use the **setbuf** subroutine to change the stream buffering strategy.

When an output stream is unbuffered, information is queued for writing on the destination file or terminal as soon as it is written. When an output stream is buffered, many characters are saved and written as a block. When an output stream is line-buffered, each line of output is queued for writing on the destination terminal as soon as the line is completed (that is, as soon as a new-line character is written or terminal input is requested).

## Parameters

<i>Stream</i>	Points to the file structure of an open file.
<i>Character</i>	Specifies a character to be written.
<i>Word</i>	Specifies a word to be written (not portable because word length and byte-ordering are machine-dependent).

## Return Values

Upon successful completion, these functions each return the value written. If these functions fail, they return the constant **EOF**. They fail if the *Stream* parameter is not open for writing, or if the output file size cannot be increased. Because the **EOF** value is a valid integer, you should use the **ferror** subroutine to detect **putc** errors.

## Error Codes

The **fputc** subroutine will fail if either the *Stream* is unbuffered or the *Stream* buffer needs to be flushed, and:

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> flag is set for the file descriptor underlying <i>Stream</i> and the process would be delayed in the write operation.
<b>EBADF</b>	The file descriptor underlying <i>Stream</i> is not a valid file descriptor open for writing.
<b>EFBIG</b>	An attempt was made to write a file that exceeds the file size of the process limit or the maximum file size.
<b>EFBIG</b>	The file is a regular file and an attempt was made to write at or beyond the offset maximum.
<b>EINTR</b>	The write operation was terminated due to the receipt of a signal, and either no data was transferred or the implementation does not report partial transfers for this file. <b>Note:</b> Depending upon which library routine the application binds to, this subroutine may return <b>EINTR</b> . Refer to the <b>signal</b> Subroutine regarding <b>sa_restart</b> .
<b>EIO</b>	A physical I/O error has occurred, or the process is a member of a background process group attempting to perform a <b>write</b> subroutine to its controlling terminal, the <b>TOSTOP</b> flag is set, the process is neither ignoring nor blocking the <b>SIGTTOU</b> signal and the process group of the process is orphaned. This error may also be returned under implementation-dependent conditions.
<b>ENOSPC</b>	There was no free space remaining on the device containing the file.
<b>EPIPE</b>	An attempt is made to write to a pipe or first-in-first-out (FIFO) that is not open for reading by any process. A <b>SIGPIPE</b> signal will also be sent to the process.

The **fputc** subroutine may fail if:

<b>ENOMEM</b>	Insufficient storage space is available.
<b>ENXIO</b>	A request was made of a nonexistent device, or the request was outside the capabilities of the device.

## Related Information

The **fclose** or **fflush** (“fclose or fflush Subroutine” on page 253) subroutine, **feof**, **ferror**, **clearerr**, or **fileno** (“feof, ferror, clearerr, or fileno Macro” on page 270) subroutine, **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fread** or **fwrite** (“fread or fwrite Subroutine” on page 310) subroutine, **getc**, **fgetc**, **getchar**, or **getw** (“getc, getchar, fgetc, or getw Subroutine” on page 347) subroutine, **getwc**, **fgetwc**, or **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **printf**, **fprintf**, **sprintf**, **NLprintf**, **NLfprintf**, **NLsprintf**, or **wsprintf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **putwc**, **fputwc**, or **putwchar** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **puts** or **fputs** (“puts or fputs Subroutine” on page 1331) subroutine, **setbuf** subroutine.

---

## putconfattr Subroutine

### Purpose

Accesses system information in the system information database.

### Library

Security Library (**libc.a**)

## Syntax

```
#include <usersec.h>
#include <userconf.h>
```

```
int putconfattrs (Table, Attributes, Count)
char * Table;
dbattr_t * Attributes;
int Count
```

## Description

The **putconfattrs** subroutine writes one or more attributes into the system information database. If the database is not already open, the subroutine does an implicit open for reading and writing. Data changed by **putconfattrs** must be explicitly committed by calling the **putconfattr** subroutine with a *Type* parameter specifying the **SEC\_COMMIT** value. Until the data is committed, only **get** subroutine calls within the process return the written data.

The *Attributes* array contains information about each attribute that is to be written. The **dbattr\_t** data structure contains the following fields:

### **attr\_name**

The name of the desired attribute.

### **attr\_idx**

Used internally by the **putconfattrs** subroutine.

### **attr\_type**

The type of the desired attribute. The list of attribute types is defined in the **usersec.h** header file.

### **attr\_flag**

The results of the request to write the desired attribute.

### **attr\_un**

A union containing the values to be written. Its union members that follow correspond to the definitions of the **attr\_char**, **attr\_int**, **attr\_long**, and **attr\_llong** macros, respectively:

#### **un\_char**

Attributes of type **SEC\_CHAR** and **SEC\_LIST** store a pointer to the value to be written.

**un\_int** Attributes of type **SEC\_INT** and **SEC\_BOOL** contain the value of the attribute to be written.

#### **un\_long**

Attributes of type **SEC\_LONG** contain the value of the attribute to be written.

#### **un\_llong**

Attributes of type **SEC\_LLONG** contain the value of the attribute to be written.

### **attr\_domain**

The authentication domain containing the attribute. The **putconfattrs** subroutine stores the name of the authentication domain that was used to write this attribute if it is not initialized by the caller. The **putconfattrs** subroutine is responsible for managing the memory referenced by this pointer.

Use the **setuserdb** and **enduserdb** subroutines to open and close the system information database. Failure to explicitly open and close the system information database can result in loss of memory and performance.

## Parameters

*Table* The system information table containing the desired attributes. The list of valid system information tables is defined in the **userconf.h** header file.

<i>Attributes</i>	A pointer to an array of one or more elements of type <b>dbattr_t</b> . The list of system attributes is defined in the <b>usersec.h</b> header file.
<i>Count</i>	The number of array elements in <i>Attributes</i> .

## Security

Files accessed:

Mode	File
rw	<b>/etc/security/.ids</b>
rw	<b>/etc/security/audit/config</b>
rw	<b>/etc/security/audit/events</b>
rw	<b>/etc/security/audit/objects</b>
rw	<b>/etc/security/login.cfg</b>
rw	<b>/etc/security/portlog</b>
rw	<b>/etc/security/roles</b>
rw	<b>/usr/lib/security/methods.cfg</b>
rw	<b>/usr/lib/security/mkuser.sys</b>

## Return Values

The **putconfattr** subroutine, when successfully completed, returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **putconfattr** subroutine fails if one or more of the following are true:

<b>EACCES</b>	The system information database could not be accessed for writing.
<b>EINVAL</b>	The <i>Table</i> parameter is the NULL pointer.
<b>EINVAL</b>	The <i>Attributes</i> parameter does not point to valid data for the requested attribute. Limited testing is possible and all errors might not be detected.
<b>EINVAL</b>	The <i>Count</i> parameter is less than or equal to 0.
<b>ENOENT</b>	The specified <i>Table</i> does not exist.

If the **putconfattr** subroutine fails to write an attribute, one or more of the following errors is returned in the **attr\_flag** field of the corresponding *Attributes* element:

<b>EACCES</b>	The user does not have access to the attribute specified in the <b>attr_name</b> field.
<b>EINVAL</b>	The <b>attr_type</b> field in the <i>Attributes</i> entry contains an invalid type.
<b>EINVAL</b>	The <b>attr_un</b> field in the <i>Attributes</i> entry does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors might not be detected.
<b>ENOATTR</b>	The <b>attr_name</b> field in the <i>Attributes</i> entry specifies an attribute that is not defined for this system table.

## Related Information

The **putuserattr** (“getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454) subroutine.

The setuserdb Subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## putenv Subroutine

### Purpose

Sets an environment variable.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int putenv ( String)  
char *String;
```

### Description

**Attention:** Unpredictable results can occur if a subroutine passes the **putenv** subroutine a pointer to an automatic variable and then returns while the variable is still part of the environment.

The **putenv** subroutine sets the value of an environment variable by altering an existing variable or by creating a new one. The *String* parameter points to a string of the form *Name=Value*, where *Name* is the environment variable and *Value* is the new value for it.

The memory space pointed to by the *String* parameter becomes part of the environment, so that altering the string effectively changes part of the environment. The space is no longer used after the value of the environment variable is changed by calling the **putenv** subroutine again. Also, after the **putenv** subroutine is called, environment variables are not necessarily in alphabetical order.

The **putenv** subroutine manipulates the **environ** external variable and can be used in conjunction with the **getenv** subroutine. However, the *EnvironmentPointer* parameter, the third parameter to the main subroutine, is not changed.

The **putenv** subroutine uses the **malloc** subroutine to enlarge the environment.

### Parameters

*String*      A pointer to the *Name=Value* string.

### Return Values

Upon successful completion, a value of 0 is returned. If the **malloc** subroutine is unable to obtain sufficient space to expand the environment, then the **putenv** subroutine returns a nonzero value.

### Related Information

The **exec**: **execl**, **execv**, **execle**, **execlp**, **execvp**, or **exec** (“exec: execl, execle, execlp, execv, execve, execvp, or exec Subroutine” on page 236) subroutine, **getenv** (“getenv Subroutine” on page 365) subroutine, **malloc** (“malloc, free, realloc, calloc, mallopt, mallinfo, mallinfo\_heap, alloca, valloc, or posix\_memalign Subroutine” on page 778) subroutine.

---

## putgrent Subroutine

### Purpose

Updates group descriptions.

## Library

Standard C Library (**libc.a**)

## Syntax

```
int putgrent (grp, fp)
struct group *grp;
FILE *fp;
```

## Description

The **putgrent** subroutine updates group descriptions. The *grp* parameter is a pointer to a group structure, as created by the **getgrent**, **getgrgid**, and **getgrnam** subroutines.

The **putgrent** subroutine writes a line on the stream specified by the *fp* parameter. The stream matches the format of **/etc/group**.

The **gr\_passwd** field of the line written is always set to ! (exclamation point).

## Parameters

<i>grp</i>	Pointer to a group structure.
<i>fp</i>	Specifies the stream to be written to.

## Return Values

The **putgrent** subroutine returns a value of 0 upon successful completion. If **putgrent** fails, a nonzero value is returned.

## Files

**/etc/group**

**/etc/security/group**

## Related Information

The “getgrent, getgrgid, getgrnam, setgrent, or endgrent Subroutine” on page 371.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## putgroupattrs Subroutine

### Purpose

Stores multiple group attributes in the group database.

### Library

Security Library (**libsec.a**)

### Syntax

```
#include <usersec.h>
```

```
int putgroupattrs (Group, Attributes, Count)
char * Group;
dbattr_t * Attributes;
int Count
```

## Description

The **putgroupattrs** subroutine writes multiple group attributes into the group database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by **putgroupattrs** must be explicitly committed by calling the **putgroupattr** subroutine with a *Type* parameter specifying the **SEC\_COMMIT** value. Until the data is committed, only **get** subroutine calls within the process return the written data.

The *Attributes* array contains information about each attribute that is to be written. Each element in the *Attributes* array must be examined upon a successful call to **putgroupattrs** to determine if the *Attributes* array entry was successfully put. The **dbattr\_t** data structure contains the following fields:

### attr\_name

The name of the desired attribute.

### attr\_idx

Used internally by the **putgroupattrs** subroutine.

### attr\_type

The type of the desired attribute. The list of attribute types is defined in the **usersec.h** header file.

### attr\_flag

The results of the request to write the desired attribute.

### attr\_un

A union containing the values to be written. Its union members that follow correspond to the definitions of the **attr\_char**, **attr\_int**, **attr\_long**, and **attr\_llong** macros, respectively:

#### un\_char

Attributes of type **SEC\_CHAR** and **SEC\_LIST** store a pointer to the value to be written.

**un\_int** Attributes of type **SEC\_INT** and **SEC\_BOOL** contain the value of the attribute to be written.

#### un\_long

Attributes of type **SEC\_LONG** contain the value of the attribute to be written.

#### un\_llong

Attributes of type **SEC\_LLONG** contain the value of the attribute to be written.

### attr\_domain

The authentication domain containing the attribute. The **putgroupattrs** subroutine stores the name of the authentication domain that was used to write this attribute if it is not initialized by the caller. The **putgroupattrs** subroutine is responsible for managing the memory referenced by this pointer. If **attr\_domain** is specified for an attribute, the put request is sent only to that domain.

If **attr\_domain** is not specified (that is, set to NULL), **putgroupattrs** attempts to put the attributes to the first domain associated with the user. All put requests for the attributes with a NULL **attr\_domain** are sent to the same domain. In other words, values cannot be put into different domains where **attr\_domain** is unspecified; **attr\_domain** is set to the name of the domain where the value is put and returned to the invoker.

When **attr\_domain** is not specified, the list of searchable domains can be restricted to a particular domain by using the **setauthdb** function call.

Use the **setuserdb** and **enduserdb** subroutines to open and close the group database. Failure to explicitly open and close the group database can result in loss of memory and performance.



## Parameters

<i>Group</i>	Specifies the name of the group for which the attributes are to be written.
<i>Attributes</i>	A pointer to an array of one or more elements of type <b>dbattr_t</b> . The list of group attributes is defined in the <b>usersec.h</b> header file.
<i>Count</i>	The number of array elements in <i>Attributes</i> .

## Security

Files accessed:

Mode	File
rw	/etc/group
rw	/etc/security/group
rw	/etc/security/smitacl.group

## Return Values

The **putgroupattrs** subroutine returns a value of 0 if the *Group* exists, even in the case when no attributes in the *Attributes* array were successfully updated. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **putgroupattrs** subroutine fails if one or more of the following are true:

<b>EACCES</b>	The system information database could not be accessed for writing.
<b>EINVAL</b>	The <i>Group</i> parameter is the NULL pointer.
<b>EINVAL</b>	The <i>Attributes</i> parameter does not point to valid data for the requested attribute. Limited testing is possible and all errors might not be detected.
<b>EINVAL</b>	The <i>Count</i> parameter is less than or equal to 0.
<b>ENOENT</b>	The specified <i>Group</i> does not exist.

If the **putgroupattrs** subroutine fails to write an attribute, one or more of the following errors is returned in the **attr\_flag** field of the corresponding *Attributes* element:

<b>EACCES</b>	The user does not have access to the attribute specified in the <b>attr_name</b> field.
<b>EINVAL</b>	The <b>attr_type</b> field in the <i>Attributes</i> entry contains an invalid type.
<b>EINVAL</b>	The <b>attr_un</b> field in the <i>Attributes</i> entry does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors might not be detected.
<b>ENOATTR</b>	The <b>attr_name</b> field in the <i>Attributes</i> entry specifies an attribute that is not defined for this group.

## Examples

The following sample test program displays the output to a call to **putgroupattrs**. In this example, the system has a user named foo and a group named bar.

```
#include <stdio.h>
#include <strings.h>
#include <string.h>
#include <usersec.h>
```

```
char * CommaToNSL(char *);
```

```
#define NATTR 2 /* Number of attributes to be put. */
#define GROUPNAME "bar" /* Group name. */
```

```

#define DOMAIN "files" /* Domain where attributes are going to put. */

main(int argc, char *argv[]) {
    int rc;
    int i;
    dbattr_t attributes[NATTR];

        /* Open the group database */
        setuserdb(S_WRITE);

    /* Valid put */

    attributes[0].attr_name = S_ADMIN;
        attributes[0].attr_type = SEC_BOOL;
    attributes[0].attr_domain = DOMAIN;
    attributes[0].attr_char = strdup("false");

    /* Valid put */

    attributes[1].attr_name = S_USERS;
    attributes[1].attr_type = SEC_LIST;
    attributes[1].attr_domain = DOMAIN;
    attributes[1].attr_char = CommaToNSL("foo");

    rc = putgroupattrs(GROUPNAME, attributes, NATTR);

    if (rc) {
        printf("putgroupattrs failed \n");
        goto clean_exit;
    }

    for (i = 0; i < NATTR; i++) {
        if (attributes[i].attr_flag)
            printf("Put failed for attribute %s. errno = %d \n",
                attributes[i].attr_name, attributes[i].attr_flag);
        else
            printf("Put succeeded for attribute %s \n",
                attributes[i].attr_name);
    }

clean_exit:
    enduserdb();

    if (attributes[0].attr_char)
        free(attributes[0].attr_char);

        if (attributes[1].attr_char)
            free(attributes[1].attr_char);

    exit(rc);
}

/*
 * Returns a new NSL created from a comma separated list.
 * The comma separated list is unmodified.
 */
char *
CommaToNSL(char *CommaList)
{
    char *NSL = (char *) NULL;
    char *s;

    if (!CommaList)
        return(NSL);
}

```

```

    if (!(NSL = (char *) malloc(strlen(CommaList) + 2)))
        return(NSL);

    strcpy(NSL, CommaList);

    for (s = NSL; *s; s++)
        if (*s == ',')
            *s = '\\0';

    *(++s) = '\\0';
}

```

The following output for the call is expected:

```

Put succeeded for attribute admin
Put succeeded for attribute users

```

## Related Information

The setuserdb Subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

## puts or fputs Subroutine

### Purpose

Writes a string to a stream.

### Library

Standard I/O Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
int puts (String)
const char *String;
```

```
int fputs (String, Stream)
const char *String;
FILE *Stream;
```

### Description

The **puts** subroutine writes the string pointed to by the *String* parameter to the standard output stream, **stdout**, and appends a new-line character to the output.

The **fputs** subroutine writes the null-terminated string pointed to by the *String* parameter to the output stream specified by the *Stream* parameter. The **fputs** subroutine does not append a new-line character.

Neither subroutine writes the terminating null character.

After the **fputwc**, **putwc**, **fputc**, **fputs**, **puts**, or **putw** subroutine runs successfully, and before the next successful completion of a call either to the **fflush** or **fclose** subroutine on the same stream or a call to the **exit** or **abort** subroutine, the `st_ctime` and `st_mtime` fields of the file are marked for update.

## Parameters

*String* Points to a string to be written to output.  
*Stream* Points to the **FILE** structure of an open file.

## Return Values

Upon successful completion, the **puts** and **fputs** subroutines return the number of characters written. Otherwise, both subroutines return **EOF**, set an error indicator for the stream and set the **errno** global variable to indicate the error. This happens if the routines try to write to a file that has not been opened for writing.

## Error Codes

If the **puts** or **fputs** subroutine is unsuccessful because the output stream specified by the *Stream* parameter is unbuffered or the buffer needs to be flushed, it returns one or more of the following error codes:

**EAGAIN** Indicates that the **O\_NONBLOCK** flag is set for the file descriptor specified by the *Stream* parameter and the process would be delayed in the write operation.

**EBADF** Indicates that the file descriptor specified by the *Stream* parameter is not a valid file descriptor open for writing.

**EFBIG** Indicates that an attempt was made to write to a file that exceeds the process' file size limit or the systemwide maximum file size.

**EINTR** Indicates that the write operation was terminated due to receipt of a signal and no data was transferred. **Note:** Depending upon which library routine the application binds to, this subroutine may return **EINTR**. Refer to the **signal** subroutine regarding the **SA\_RESTART** bit.

**EIO** Indicates that the process is a member of a background process group attempting to perform a write to its controlling terminal, the **TOSTOP** flag is set, the process is neither ignoring or blocking the **SIGTTOU** signal, and the process group of the process has no parent process.

**ENOSPC** Indicates that there was no free space remaining on the device containing the file specified by the *Stream* parameter.

**EPIPE** Indicates that an attempt is made to write to a pipe or first-in-first-out (FIFO) that is not open for reading by any process. A **SIGPIPE** signal will also be sent to the process.

**ENOMEM** Indicates that insufficient storage space is available.

**ENXIO** Indicates that a request was made of a nonexistent device, or the request was outside the capabilities of the device.

## Related Information

The **fopen**, **freopen**, or **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fread**, or **fwrite** (“fread or fwrite Subroutine” on page 310) subroutine, **gets** or **fgets** (“gets or fgets Subroutine” on page 434) subroutine, **getws** or **fgetws** (“getws or fgetws Subroutine” on page 480) subroutine, **printf**, **fprintf**, and **sprintf** (“printf, fprintf, sprintf, snprintf, vsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **putc**, **putchar**, **fputc**, or **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **putwc**, **putwchar**, or **fputwc** (“putwc, putwchar, or fputwc Subroutine” on page 1339) subroutine, **putws** or **fputws** (“putws or fputws Subroutine” on page 1341) subroutine.

The **feof**, **ferror**, **clearerr**, or **fileno** (“feof, ferror, clearerr, or fileno Macro” on page 270) macros.

List of String Manipulation Services.

Subroutines Overview in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## putuserattrs Subroutine

### Purpose

Stores multiple user attributes in the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int putuserattrs (User, Attributes, Count)
char * User;
dbattr_t * Attributes;
int Count
```

### Description

The **putuserattrs** subroutine writes multiple user attributes into the user database. If the database is not already open, this subroutine does an implicit open for reading and writing. Data changed by **putuserattrs** must be explicitly committed by calling the **putuserattr** subroutine with a *Type* parameter specifying the **SEC\_COMMIT** value. Until the data is committed, only **get** subroutine calls within the process return the written data.

The *Attributes* array contains information about each attribute that is to be written. Each element in the *Attributes* array must be examined upon a successful call to **putuserattrs** to determine if the *Attributes* array entry was successfully put. Please see **putuserattr** man page for the supported attributes. The **dbattr\_t** data structure contains the following fields:

#### attr\_name

The name of the desired attribute.

#### attr\_idx

Used internally by the **putuserattrs** subroutine.

#### attr\_type

The type of the desired attribute. The list of attribute types is defined in the **usersec.h** header file.

#### attr\_flag

The results of the request to write the desired attribute.

#### attr\_un

A union containing the returned values. Its union members that follow correspond to the definitions of the **attr\_char**, **attr\_int**, **attr\_long**, and **attr\_llong** macros, respectively:

##### un\_char

Attributes of type **SEC\_CHAR** and **SEC\_LIST** contain a pointer to the value to be written.

**un\_int** Attributes of type **SEC\_INT** and **SEC\_BOOL** contain the value of the attribute to be written.

##### un\_long

Attributes of type **SEC\_LONG** contain the value of the attribute to be written.

##### un\_llong

Attributes of type **SEC\_LLONG** contain the value of the attribute to be written.

#### attr\_domain

The authentication domain containing the attribute. The **putuserattrs** subroutine stores the name

of the authentication domain that was used to write this attribute if it is not initialized by the caller. The **putuserattrs** subroutine is responsible for managing the memory referenced by this pointer. If **attr\_domain** is specified for an attribute, the put request is sent only to that domain. If **attr\_domain** is not specified (that is, set to NULL), **putuserattrs** attempts to put the attributes to the first domain associated with the user. All put requests for the attributes with a NULL **attr\_domain** are sent to the same domain. In other words, values cannot be put into different domains where **attr\_domain** is unspecified; **attr\_domain** is set to the name of the domain where the value is put and returned to the invoker. When **attr\_domain** is not specified, the list of searchable domains can be restricted to a particular domain by using the **setauthdb** function call.

Use the **setuserdb** and **enduserdb** subroutines to open and close the user database. Failure to explicitly open and close the user database can result in loss of memory and performance.

## Parameters

<i>User</i>	Specifies the name of the user for which the attributes are to be written.
<i>Attributes</i>	A pointer to an array of one or more elements of type <b>dbattr_t</b> . The list of user attributes is defined in the <b>usersec.h</b> header file.
<i>Count</i>	The number of array elements in <i>Attributes</i> .

## Security

Files accessed:

Mode	File
rw	<b>/etc/group</b>
rw	<b>/etc/passwd</b>
rw	<b>/etc/security/audit/config</b>
rw	<b>/etc/security/environ</b>
rw	<b>/etc/security/group</b>
rw	<b>/etc/security/lastlog</b>
rw	<b>/etc/security/limits</b>
rw	<b>/etc/security/passwd</b>
rw	<b>/etc/security/pwdhist.dir</b>
rw	<b>/etc/security/pwdhist.pag</b>
rw	<b>/etc/security/smitacl.user</b>
rw	<b>/etc/security/user.roles</b>

## Return Values

The **putuserattrs** subroutine returns a value of 0 if the *User* exists, even in the case when no attributes in the *Attributes* array were successfully updated. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **putuserattrs** subroutine fails if one or more of the following is true:

<b>EACCES</b>	The system information database could not be accessed for writing.
<b>EINVAL</b>	The <i>User</i> parameter is the NULL pointer.
<b>EINVAL</b>	The <i>Attributes</i> parameter does not point to valid data for the requested attribute. Limited testing is possible and all errors might not be detected.
<b>EINVAL</b>	The <i>Attributes</i> parameter does not point to valid data for the requested attribute. Limited testing is possible and all errors might not be detected.
<b>ENOENT</b>	The specified <i>User</i> parameter does not exist.

If the **putuserattrs** subroutine fails to write an attribute, one or more of the following errors is returned in the **attr\_flag** field of the corresponding *Attributes* element:

<b>EACCES</b>	The user does not have access to the attribute specified in the <i>attr_name</i> field.
<b>EINVAL</b>	The <b>attr_type</b> field in the <i>Attributes</i> entry contains an invalid type.
<b>EINVAL</b>	The <b>attr_un</b> field in the <i>Attributes</i> entry does not point to a valid buffer or to valid data for this type of attribute. Limited testing is possible and all errors might not be detected.
<b>ENOATTR</b>	The <b>attr_name</b> field in the <i>Attributes</i> entry specifies an attribute that is not defined for this user.

## Examples

The following sample test program displays the output to a call to **putuserattrs**. In this example, the system has a user named foo.

```
#include <stdio.h>
#include <strings.h>
#include <string.h>
#include <usersec.h>

char * CommaToNSL(char *);

#define NATTR 4 /* Number of attributes to be put */
#define USERNAME "foo" /* User name */
#define DOMAIN "files" /* domain where attributes are going to put. */

main(int argc, char *argv[]) {
    int rc;
    int i;
    dbattr_t attributes[NATTR];

    /* Open the user database */
    setuserdb(S_WRITE);

    /* Valid put */

    attributes[0].attr_name = S_GECOS;
    attributes[0].attr_type = SEC_CHAR;
    attributes[0].attr_domain = DOMAIN;
    attributes[0].attr_char = strdup("I am foo");

    /* Invalid put */

    attributes[1].attr_name = S_LOGINCHK;
    attributes[1].attr_type = SEC_BOOL;
    attributes[1].attr_domain = DOMAIN;
    attributes[1].attr_char = strdup("allow");

    /* Valid put */

    attributes[2].attr_name = S_MAXAGE;
    attributes[2].attr_type = SEC_INT;
    attributes[2].attr_domain = DOMAIN;
    attributes[2].attr_int = 10;

    /* Valid put */

    attributes[3].attr_name = S_GROUPS;
    attributes[3].attr_type = SEC_LIST;
    attributes[3].attr_domain = DOMAIN;
    attributes[3].attr_char = CommaToNSL("staff,system");

    rc = putuserattrs(USERNAME, attributes, NATTR);
```

```

if (rc) {
    printf("putuserattrs failed \n");
    goto clean_exit;
}

for (i = 0; i < NATTR; i++) {
    if (attributes[i].attr_flag)
        printf("Put failed for attribute %s. errno = %d \n",
            attributes[i].attr_name, attributes[i].attr_flag);
    else
        printf("Put succeeded for attribute %s \n",
            attributes[i].attr_name);
}

clean_exit:
    enduserdb();

    if (attributes[0].attr_char)
        free(attributes[0].attr_char);

        if (attributes[1].attr_char)
            free(attributes[1].attr_char);

            if (attributes[3].attr_char)
                free(attributes[3].attr_char);

    exit(rc);
}

/*
 * Returns a new NSL created from a comma separated list.
 * The comma separated list is unmodified.
 */
char *
CommaToNSL(char *CommaList)
{
    char    *NSL = (char *) NULL;
    char    *s;

    if (!CommaList)
        return(NSL);

    if (!(NSL = (char *) malloc(strlen(CommaList) + 2)))
        return(NSL);

    strcpy(NSL, CommaList);

    for (s = NSL; *s; s++)
        if (*s == ',')
            *s = '\\0';

    *(++s) = '\\0';
}

```

The following output for the call is expected:

```

Put succeeded for attribute gecost
Put failed for attribute login (errno = 22)
Put succeeded for attribute maxage
Put succeeded for attribute groups

```

## Related Information

The setuserdb Subroutine.



## putuserpwx Subroutine

### Purpose

Accesses the user authentication data.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <userpw.h>
```

```
int putuserpwx (Password)  
struct userpwx *Password;
```

### Description

The **putuserpwx** subroutine modifies user authentication information. It can be used with those administrative domains that support modifying the user's encrypted password with the **putuserattrs** subroutine. The **chpassx** subroutine must be used to modify authentication information for administrative domains that do not support that functionality.

The **putuserpwx** subroutine updates or creates password authentication data for the user defined in the *Password* parameter in the administrative domain that is specified. The password entry created by the **putuserpwx** subroutine is used only if there is an ! (exclamation point) in the user's password (**S\_PWD**) attribute. The user application can use the **putuserattrs** subroutine to add an ! to this field.

The **putuserpwx** subroutine opens the authentication database read-write if no other access has taken place, but the program should call **setpwdb** (**S\_READ** | **S\_WRITE**) before calling the **putuserpwx** subroutine and **endpwdb** when access to the authentication information is no longer required.

The administrative domain specified in the **upw\_authdb** field is set by the **getuserpwx** subroutine. It must be specified by the application program if the **getuserpwx** subroutine is not used to produce the *Password* parameter.

## Parameters

### *Password*

Specifies the password structure used to update the password information for this user. The fields in a **userpw** structure are defined in the **userpw.h** file and contains the following members:

#### **upw\_name**

Specifies the user's name.

#### **upw\_passwd**

Specifies the user's encrypted password.

#### **upw\_lastupdate**

Specifies the time, in seconds, since the epoch (that is, 00:00:00 GMT, 1 January 1970), when the password was last updated.

#### **upw\_flags**

Specifies attributes of the password. This member is a bit mask of one or more of the following values, defined in the **userpw.h** file:

##### **PW\_NOCHECK**

Specifies that new passwords need not meet password restrictions in effect for the system.

##### **PW\_ADMCHG**

Specifies that the password was last set by an administrator and must be changed at the next successful use of the login or **su** command.

##### **PW\_ADMIN**

Specifies that password information for this user can only be changed by the root user.

#### **upw\_authdb**

Specifies the administrative domain containing the authentication data.

## Security

Files accessed:

Mode	File
rw	<b>/etc/security/passwd</b>

## Return Values

If successful, the **putuserpw** subroutine returns a value of 0. If the subroutine failed to update or create the password information, the **putuserpw** subroutine returns a nonzero value.

## Error Codes

The **getuserpw** subroutine fails if the following value is true:

**ENOENT**                      The user does not have an entry in the **/etc/security/passwd** file.

Subroutines invoked by the **putuserpw** subroutine can also set errors.

## Files

**/etc/security/passwd**        Contains user passwords.

## Related Information

The “getuserattr, IDtouser, nextuser, or putuserattr Subroutine” on page 454, “putgroupattrs Subroutine” on page 1327, “putuserattrs Subroutine” on page 1333, setpwdb Subroutinesetuserdb Subroutine.

---

## putwc, putwchar, or fputwc Subroutine

### Purpose

Writes a character or a word to a stream.

### Library

Standard I/O Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
wint_t putwc( Character, Stream)
```

```
wint_t Character;
```

```
FILE *Stream;
```

```
wint_t putwchar(Character)
```

```
wint_t Character;
```

```
wint_t fputwc(Character, Stream)
```

```
wint_t Character;
```

```
FILE Stream;
```

### Description

The **putwc** subroutine writes the wide character specified by the *Character* parameter to the output stream pointed to by the *Stream* parameter. The wide character is written as a multibyte character at the associated file position indicator for the stream, if defined. The subroutine then advances the indicator. If the file cannot support positioning requests, or if the stream was opened with append mode, the character is appended to the output stream.

The **putwchar** subroutine works like the **putwc** subroutine, except that **putwchar** writes the specified wide character to the standard output.

The **fputwc** subroutine works the same as the **putwc** subroutine.

Output streams, with the exception of **stderr**, are buffered by default if they refer to files, or line-buffered if they refer to terminals. The standard error output stream, **stderr**, is unbuffered by default, but using the **freopen** subroutine causes it to become buffered or line-buffered. Use the **setbuf** subroutine to change the stream's buffering strategy.

After the **fputwc**, **putwc**, **fputc**, **putc**, **fputs**, **puts**, or **putw** subroutine runs successfully, and before the next successful completion of a call either to the **fflush** or **fclose** subroutine on the same stream or to the **exit** or **abort** subroutine, the *st\_ctime* and *st\_mtime* fields of the file are marked for update.

### Parameters

<i>Character</i>	Specifies a wide character of type <b>wint_t</b> .
<i>Stream</i>	Specifies a stream of output data.

## Return Values

Upon successful completion, the **putwc**, **putwchar**, and **fputwc** subroutines return the wide character that is written. Otherwise **WEOF** is returned, the error indicator for the stream is set, and the **errno** global variable is set to indicate the error.

## Error Codes

If the **putwc**, **putwchar**, or **fputwc** subroutine fails because the stream is not buffered or data in the buffer needs to be written, it returns one or more of the following error codes:

<b>EAGAIN</b>	Indicates that the <b>O_NONBLOCK</b> flag is set for the file descriptor underlying the <i>Stream</i> parameter, delaying the process during the write operation.
<b>EBADF</b>	Indicates that the file descriptor underlying the <i>Stream</i> parameter is not valid and cannot be updated during the write operation.
<b>EFBIG</b>	Indicates that the process attempted to write to a file that already equals or exceeds the file-size limit for the process. The file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.
<b>EILSEQ</b>	Indicates that the wide-character code does not correspond to a valid character.
<b>EINTR</b>	Indicates that the process has received a signal that terminates the read operation.
<b>EIO</b>	Indicates that the process is in a background process group attempting to perform a write operation to its controlling terminal. The <b>TOSTOP</b> flag is set, the process is not ignoring or blocking the <b>SIGTTOU</b> flag, and the process group of the process is orphaned.
<b>ENOMEM</b>	Insufficient storage space is available.
<b>ENOSPC</b>	Indicates that no free space remains on the device containing the file.
<b>ENXIO</b>	Indicates a request was made of a non-existent device, or the request was outside the capabilities of the device.
<b>EPIPE</b>	Indicates that the process has attempted to write to a pipe or first-in-first-out (FIFO) that is not open for reading. The process will also receive a <b>SIGPIPE</b> signal.

## Related Information

Other wide character I/O subroutines: **fgetwc** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **fgetws** (“getws or fgetws Subroutine” on page 480) subroutine, **fputws** (“putws or fputws Subroutine” on page 1341) subroutine, **getwc** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **getwchar** (“getwc, fgetwc, or getwchar Subroutine” on page 477) subroutine, **getws** (“getws or fgetws Subroutine” on page 480) subroutine, **putws** (“putws or fputws Subroutine” on page 1341) subroutine, **ungetwc** subroutine.

Related standard I/O subroutines: **fdopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fgets** (“gets or fgets Subroutine” on page 434) subroutine, **fopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fprintf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **fputc** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **fputs** (“puts or fputs Subroutine” on page 1331) subroutine, **fread** (“fread or fwrite Subroutine” on page 310) subroutine, **freopen** (“fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288) subroutine, **fwrite** (“read or fwrite Subroutine” on page 310) subroutine, **gets** (“gets or fgets Subroutine” on page 434) subroutine, **printf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine, **putc** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **putchar** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **puts** (“puts or fputs Subroutine” on page 1331) subroutine, **putw** (“putc, putchar, fputc, or putw Subroutine” on page 1321) subroutine, **sprintf** (“printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## putws or fputws Subroutine

### Purpose

Writes a wide-character string to a stream.

### Library

Standard I/O Library (**libc.a**)

### Syntax

```
#include <stdio.h>
```

```
int putws ( String)  
const wchar_t *String;
```

```
int fputws (String, Stream)  
const wchar_t *String;  
FILE *Stream;
```

### Description

The **putws** subroutine writes the **const wchar\_t** string pointed to by the *String* parameter to the standard output stream (**stdout**) as a multibyte character string and appends a new-line character to the output. In all other respects, the **putws** subroutine functions like the **puts** subroutine.

The **fputws** subroutine writes the **const wchar\_t** string pointed to by the *String* parameter to the output stream as a multibyte character string. In all other respects, the **fputws** subroutine functions like the **fputs** subroutine.

After the **putws** or **fputws** subroutine runs successfully, and before the next successful completion of a call to the **fflush** or **fclose** subroutine on the same stream or a call to the **exit** or **abort** subroutine, the *st\_ctime* and *st\_mtime* fields of the file are marked for update.

### Parameters

*String*        Points to a string to be written to output.  
*Stream*        Points to the **FILE** structure of an open file.

### Return Values

Upon successful completion, the **putws** and **fputws** subroutines return a nonnegative number. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

### Error Codes

The **putws** or **fputws** subroutine is unsuccessful if the stream is not buffered or data in the buffer needs to be written, and one of the following errors occur:

**EAGAIN**     The **O\_NONBLOCK** flag is set for the file descriptor underlying the *Stream* parameter, which delays the process during the write operation.  
**EBADF**     The file descriptor underlying the *Stream* parameter is not valid and cannot be updated during the write operation.  
**EFBIG**     The process attempted to write to a file that already equals or exceeds the file-size limit for the process.  
**EINTR**     The process has received a signal that terminates the read operation.

<b>EIO</b>	The process is in a background process group attempting to perform a write operation to its controlling terminal. The <b>TOSTOP</b> flag is set, the process is not ignoring or blocking the <b>SIGTTOU</b> flag, and the process group of the process is orphaned.
<b>ENOSPC</b>	No free space remains on the device containing the file.
<b>EPIPE</b>	The process has attempted to write to a pipe or first-in-first-out (FIFO) that is not open for reading. The process also receives a <b>SIGPIPE</b> signal.
<b>EILSEQ</b>	The <b>wc</b> wide-character code does not correspond to a valid character.

## Related Information

Other wide-character I/O subroutines: “getwc, fgetwc, or getwchar Subroutine” on page 477, “getws or fgetws Subroutine” on page 480, “putwc, putwchar, or fputwc Subroutine” on page 1339, and **ungetwc** subroutine.

Related standard I/O subroutines: “fopen, fopen64, freopen, freopen64 or fdopen Subroutine” on page 288, “gets or fgets Subroutine” on page 434, “printf, fprintf, sprintf, snprintf, wsprintf, vprintf, vfprintf, vsprintf, or vwsprintf Subroutine” on page 1168, “putc, putchar, fputc, or putw Subroutine” on page 1321, “puts or fputs Subroutine” on page 1331, “fread or fwrite Subroutine” on page 310.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.3 National Language Support Guide and Reference*.

---

## pwdrestrict\_method Subroutine

### Purpose

Defines loadable password restriction methods.

### Library

### Syntax

```
int pwdrestrict_method (UserName, NewPassword, OldPassword, Message)
char * UserName;
char * NewPassword;
char * OldPassword;
char ** Message;
```

### Description

The **pwdrestrict\_method** subroutine extends the capability of the password restrictions software and lets an administrator enforce password restrictions that are not provided by the system software.

Whenever users change their passwords, the system software scans the **pwdchecks** attribute defined for that user for site specific restrictions. Since this attribute field can contain load module file names, for example, methods, it is possible for the administrator to write and install code that enforces site specific password restrictions.

The system evaluates the **pwdchecks** attribute’s value field in a left to right order. For each method that the system encounters, the system loads and invokes that method. The system uses the **load** subroutine to load methods. It invokes the **load** subroutine with a *Flags* value of **1** and a *LibraryPath* value of **/usr/lib**. Once the method is loaded, the system invokes the method.

To create a loadable module, use the **-e** flag of the **ld** command. Note that the name **pwdrestrict\_method** given in the syntax is a generic name. The actual subroutine name can be anything (within the compiler's name space) except **main**. What is important is, that for whatever name you choose, you must inform the **ld** command of the name so that the **load** subroutine uses that name as the entry point into the module. In the following example, the C compiler compiles the **pwdrestrict.c** file and pass **-e pwdrestrict\_method** to the **ld** command to create the method called **pwdrestrict**:

```
cc -e pwdrestrict_method -o pwdrestrict pwdrestrict.c
```

The convention of all password restriction methods is to pass back messages to the invoking subroutine. Do not print messages to **stdout** or **stderr**. This feature allows the password restrictions software to work across network connections where **stdout** and **stderr** are not valid. Note that messages must be returned in dynamically allocated memory to the invoking program. The invoking program will deallocate the memory once it is done with the memory.

There are many caveats that go along with loadable subroutine modules:

1. The values for *NewPassword* and *OldPassword* are the actual clear text passwords typed in by the user. If you copy these passwords into other parts of memory, clear those memory locations before returning back to the invoking program. This helps to prevent clear text passwords from showing up in core dumps. Also, do not copy these passwords into a file or anywhere else that another program can access. Clear text passwords should never exist outside of the process space.
2. Do not modify the current settings of the process' signal handlers.
3. Do not call any functions that will terminate the execution of the program (for example, the **exit** subroutine, the **exec** subroutine). Always return to the invoking program.
4. The code must be thread-safe.
5. The actual load module must be kept in a write protected environment. The load module and directory should be writable only by the root user.

One last note, all standard password restrictions are performed before any of the site specific methods are invoked. Thus, methods are the last restrictions to be enforced by the system.

## Parameters

<i>UserName</i>	Specifies a "local" user name.
<i>NewPassword</i>	Specifies the new password in clear text (not encrypted). This value may be a NULL pointer. Clear text passwords are always in 7 bit ASCII.
<i>OldPassword</i>	Specifies the current password in clear text (not encrypted). This value may be a NULL pointer. Clear text passwords are always in 7 bit ASCII.
<i>Message</i>	Specifies the address of a pointer to <b>malloc</b> 'ed memory containing an NLS error message. The method is expected to supply the <b>malloc</b> 'ed memory and the message.

## Return Values

The method is expected to return the following values. The return values are listed in order of precedence.

- 1 Internal error. The method could not perform its password evaluation. The method must set the **errno** variable. The method must supply an error message in *Message* unless it can't allocate memory for the message. If it cannot allocate memory, then it must return the NULL pointer in *Message*.
- 1 Failure. The password change did not meet the requirements of the restriction. The password restriction was properly evaluated and the password change was not accepted. The method must supply an error message in *Message*. The **errno** variable is ignored. Note that composition failures are cumulative, thus, even though a failure condition is returned, trailing composition methods will be invoked.
- 0 Success. The password change met the requirements of the restriction. If necessary, the method may supply a message in *Message*; otherwise, return the NULL pointer. The **errno** variable is ignored.





---

## Appendix A. Base Operating System Error Codes for Services That Require Path-Name Resolution

The following errors apply to any service that requires path name resolution:

<b>EACCES</b>	Search permission is denied on a component of the path prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points outside of the allocated address space of the process.
<b>EIO</b>	An I/O error occurred during the operation.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of a path name exceeded 255 characters and the process has the <b>DisallowTruncation</b> attribute (see the <b>ulimit</b> subroutine) or an entire path name exceeded 1023 characters.
<b>ENOENT</b>	A component of the path prefix does not exist.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOENT</b>	The path name is null.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>ESTALE</b>	The root or current directory of the process is located in a virtual file system that is unmounted.

---

### Related Information

List of File and Directory Manipulation Services.



---

## Appendix B. ODM Error Codes

When an ODM subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to one of the following values:

<b>ODMI_BAD_CLASSNAME</b>	The specified object class name does not match the object class name in the file. Check path name and permissions.
<b>ODMI_BAD_CLXNNAME</b>	The specified collection name does not match the collection name in the file.
<b>ODMI_BAD_CRIT</b>	The specified search criteria is incorrectly formed. Make sure the criteria contains only valid descriptor names and the search values are correct. For information on qualifying criteria, see "Understanding ODM Object Searches" in <i>AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs</i> .
<b>ODMI_BAD_LOCK</b>	Cannot set a lock on the file. Check path name and permissions.
<b>ODMI_BAD_TIMEOUT</b>	The time-out value was not valid. It must be a positive integer.
<b>ODMI_BAD_TOKEN</b>	Cannot create or open the lock file. Check path name and permissions.
<b>ODMI_CLASS_DNE</b>	The specified object class does not exist. Check path name and permissions.
<b>ODMI_CLASS_EXISTS</b>	The specified object class already exists. An object class must not exist when it is created.
<b>ODMI_CLASS_PERMS</b>	The object class cannot be opened because of the file permissions.
<b>ODMI_CLXNMAGICNO_ERR</b>	The specified collection is not a valid object class collection.
<b>ODMI_FORK</b>	Cannot fork the child process. Make sure the child process is executable and try again.
<b>ODMI_INTERNAL_ERR</b>	An internal consistency problem occurred. Make sure the object class is valid or contact the person responsible for the system.
<b>ODMI_INVALID_CLASS</b>	The specified file is not an object class.
<b>ODMI_INVALID_CLXN</b>	Either the specified collection is not a valid object class collection or the collection does not contain consistent data.
<b>ODMI_INVALID_PATH</b>	The specified path does not exist on the file system. Make sure the path is accessible.
<b>ODMI_LINK_NOT_FOUND</b>	The object class that is accessed could not be opened. Make sure the linked object class is accessible.
<b>ODMI_LOCK_BLOCKED</b>	Cannot grant the lock. Another process already has the lock.
<b>ODMI_LOCK_ENV</b>	Cannot retrieve or set the lock environment variable. Remove some environment variables and try again.
<b>ODMI_LOCK_ID</b>	The lock identifier does not refer to a valid lock. The lock identifier must be the same as what was returned from the <b>odm_lock</b> ("odm_lock Subroutine" on page 923) subroutine.
<b>ODMI_MAGICNO_ERR</b>	The class symbol does not identify a valid object class.
<b>ODMI_MALLOC_ERR</b>	Cannot allocate sufficient storage. Try again later or contact the person responsible for the system.
<b>ODMI_NO_OBJECT</b>	The specified object identifier did not refer to a valid object.
<b>ODMI_OPEN_ERR</b>	Cannot open the object class. Check path name and permissions.
<b>ODMI_OPEN_PIPE</b>	Cannot open a pipe to a child process. Make sure the child process is executable and try again.
<b>ODMI_PARAMS</b>	The parameters passed to the subroutine were not correct. Make sure there are the correct number of parameters and that they are valid.
<b>ODMI_READ_ONLY</b>	The specified object class is opened as read-only and cannot be modified.
<b>ODMI_READ_PIPE</b>	Cannot read from the pipe of the child process. Make sure the child process is executable and try again.
<b>ODMI_TOOMANYCLASSES</b>	Too many object classes have been accessed. An application can only access less than 1024 object classes.
<b>ODMI_UNLINKCLASS_ERR</b>	Cannot remove the object class from the file system. Check path name and permissions.
<b>ODMI_UNLINKCLXN_ERR</b>	Cannot remove the object class collection from the file system. Check path name and permissions.

**ODMI\_UNLOCK**

Cannot unlock the lock file. Make sure the lock file exists.

---

## **Related Information**

List of ODM Commands and Subroutines in *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*.

---

## Appendix C. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Dept. LRAS/Bldg. 003  
11400 Burnet Road  
Austin, TX 78758-3498  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

---

## Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX  
AIX 5L  
IBM  
POWER4  
POWER5  
POWER6  
PTX

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

---

# Index

## Special characters

- \_atojis macro 576
- \_check\_lock Subroutine 127
- \_clear\_lock Subroutine 128
- \_edata identifier 224
- \_end identifier 224
- \_exit subroutine 243
- \_Exit subroutine 243
- \_extext identifier 224
- \_jstoa macro 576
- \_lazySetErrorHandler Subroutine 584
- \_tolower macro 576
- \_toupper macro 576
- \_tolower subroutine 183
- \_toupper subroutine 183
- /etc/filesystems file
  - accessing entries 369
- /etc/hosts file
  - closing 911
  - retrieving host entries 910
- /etc/utmp file
  - accessing entries 474

## Numerics

- 3-byte integers
  - converting 586

## A

- a64l subroutine 1
- abort subroutine 2
- abs subroutine 3
- absinterval subroutine 386
- absolute path names
  - copying 479
  - determining 479
- absolute value subroutines
  - cabs 129
  - cabsf 129
  - cabsl 129
  - fabsf 249
- absolute values
  - computing complex 528
  - imaxabs 535
- access control attributes
  - setting 144
- access control information
  - changing 8
  - retrieving 10
  - setting 12, 14, 17, 25
- access control subroutines
  - acl\_chg 8
  - acl\_fchg 8
  - acl\_fget 10
  - acl\_fput 12
  - acl\_fset 14

access control subroutines (*continued*)

- acl\_get 10
- acl\_put 12
- acl\_set 14
- aclx\_convert 16
- aclx\_fget 17
- aclx\_fput 25
- aclx\_get 17
- aclx\_gettypeinfo 20
- aclx\_gettypes 21
- aclx\_print 23
- aclx\_printStr 23
- aclx\_put 25
- aclx\_scan 27
- aclx\_scanStr 27
- chacl 144
- chmod 148
- chown 151
- chownx 151
- fchacl 144
- fchmod 148
- fchown 151
- fchownx 151
- frevoke 313

access subroutine 4

accessx subroutine 4

accounting subroutines

- addproj 31
- addprojdb 32
- chprojattr 158
- chprojattrdb 159
- getfirstprojdb 368
- getnextprojdb 396
- getproj 418
- getprojdb 419
- getprojs 420
- proj\_execve 1178
- projdballoc 1179
- projdbfinit 1179
- projdbfree 1181

acct subroutine 7

acl\_chg subroutine 8

acl\_fchg subroutine 8

acl\_fget subroutine 10

acl\_fput subroutine 12

acl\_fset subroutine 14

acl\_get subroutine 10

acl\_put subroutine 12

acl\_set subroutine 14

aclx\_convert subroutine 16

aclx\_fget subroutine 17

aclx\_fput subroutine 25

aclx\_get subroutine 17

aclx\_gettypeinfo subroutine 20

aclx\_gettypes subroutine 21

aclx\_print subroutine 23

aclx\_printStr subroutine 23

aclx\_put subroutine 25

- aclx\_scan subroutine 27
- aclx\_scanStr subroutine 27
- acos subroutine 29
- acosf subroutine 29
- acosh subroutine 30
- acoshf subroutine 30
- acoshl subroutine 30
- acosl subroutine 29
- addproj subroutine 31
- addprojdb subroutine 32
- address identifiers 224
- addssys subroutine 33
- adjtime subroutine 35
- advance subroutine 177
- Advanced Accounting subroutines
  - agg\_arm\_stat subroutine 36
  - agg\_lpar\_stat subroutine 36
  - agg\_proc\_stat subroutine 36
  - buildproclist subroutine 125
  - buildtranlist subroutine 126
  - free\_agg\_list subroutine 36
  - freetranlist subroutine 126
  - getarmlist subroutine 414
  - getfilehdr subroutine 367
  - getparlist subroutine 414
  - getproclist subroutine 414
- agg\_arm\_stat subroutine 36
- agg\_lpar\_stat subroutine 36
- agg\_proc\_stat subroutine 36
- aio\_cancel subroutine 38
- aio\_error subroutine 42
- aio\_fsync subroutine 45
- aio\_nwait subroutine 47
- aio\_nwait\_timeout subroutine 49
- aio\_read subroutine 50
- aio\_return subroutine 55
- aio\_suspend subroutine 58
- aio\_write subroutine 61
- alarm subroutine 386
- alloca subroutine 778
- allocaimb Subroutine 68
- Application Programming Interface
  - perfstat
    - cpu 1003
    - cpu\_total 1004, 1013, 1017
    - disk\_total 1006, 1011, 1015
    - diskpath 1009
    - netbuffer 1014
    - pagingspace 1018
    - protocol 1022
    - reset 1023
- arc sine subroutines
  - asinf 91
- arc tangent subroutines
  - atan2f 93
  - atan2l 93
  - atanf 94
  - atanl 94
- archive files
  - reading headers 696
- ARM Subroutines
  - arm\_end 69
  - arm\_end Dual Call 71
  - arm\_getid 73
  - arm\_getid Dual Call 75
  - arm\_init 77
  - arm\_init Dual Call 79
  - arm\_start 81
  - arm\_start Dual Call 82
  - arm\_stop 84
  - arm\_stop Dual Call 86
  - arm\_update 88
  - arm\_update Dual Call 89
- ASCII strings
  - converting to floating-point numbers 96
  - converting to Internet addresses 557
- asctime subroutine 200
- asctime\_r subroutine 206
- asctime64 subroutine 202
- asctime64\_r subroutine 204
- asin subroutine 91
- asinf subroutine 91
- asinh subroutine 90
- asinhf subroutine 90
- asinhl subroutine 90
- asinl subroutine 91
- assert macro 92
- asynchronous I/O
  - reading 50
  - writing 61
- asynchronous I/O requests
  - canceling 38
  - listing 720
  - retrieving error status 42
  - retrieving return status 55
  - suspending 58
  - synchronizing asynchronous files 45
- atan subroutine 94
- atan2 subroutine 93
- atan2f subroutine 93
- atan2l subroutine 93
- atanf subroutine 94
- atanh subroutine 95
- atanhf subroutine 95
- atanhl subroutine 95
- atanl subroutine 94
- atexit subroutine 243
- atof subroutine 96
- atoff subroutine 96
- atojis subroutine 576
- atol subroutine 98
- atoll subroutine 98
- atomic access subroutines
  - compare\_and\_swap 176
  - fetch\_and\_add 271
  - fetch\_and\_and 272
  - fetch\_and\_or 272
- audit bin files
  - compressing and uncompressing 108
  - establishing 100



- audit records
  - generating 104
  - reading 111
  - writing 112
- audit subroutine 98
- audit trail files
  - appending records 104
- auditbin subroutine 100
- auditevents subroutine 102
- auditing modes 105
- auditing subroutines
  - audit 98
  - auditbin 100
  - auditevents 102
  - auditlog 104
  - auditobj 105
  - auditpack 108
  - auditproc 109
  - auditread 111
  - auditwrite 112
- auditlog subroutine 104
- auditobj subroutine 105
- auditpack subroutine 108
- auditproc subroutine 109
- auditread, auditread\_r subroutines 111
- auditwrite subroutine 112
- authenticate 113
- authenticatex subroutine 115
- authentication subroutines
  - ckuseracct 164
  - ckuserID 166
  - crypt 191
  - encrypt 191
  - getlogin 393
  - getpass 402
  - getuserpw 468
  - newpass 901
  - putuserpw 468
  - setkey 191
- authorizations 467
- authorizations, compare 790
- auxiliary areas
  - creating 537
  - destroying 537
  - drawing 538
  - hiding 539
  - processing 551

## B

- base 10 logarithm functions
  - log10f 743
- base 2 logarithm functions
  - log2 745
  - log2f 745
  - log2l 745
- basename Subroutine 117
- baud rates
  - getting and setting 142
- bcmp subroutine 118
- bcopy subroutine 118

- beep levels
  - setting 539
- BeginCriticalSection Subroutine 226
- Bessel functions
  - computing 119
- binary files
  - reading 310
- binary searches 123
- binding a process to a processor 120
- bit string operations 118
- box characters
  - shaping 691
- brk subroutine 122
- bsearch subroutine 123
- btowc subroutine 124
- buffered data
  - writing to streams 253
- buildproclist subroutine 125
- buildtranlist subroutine 126
- byte string operations 118
- bzero subroutine 118

## C

- cabs subroutine 129
- cabsf subroutine 129
- cabsl subroutine 129
- cacos subroutine 129
- cacosf subroutine 129
- cacosh subroutines 130
- cacoshf subroutine 130
- cacoshl subroutine 130
- cacosl subroutine 129
- calloc subroutine 778
- carg subroutine 131
- cargf subroutine 131
- cargl subroutine 131
- casin subroutine 131
- casinf subroutine 131
- casinfh subroutine 132
- casinh subroutines 132
- casinl subroutine 131
- casinlh subroutine 132
- catan subroutine 132
- catanf subroutine 132
- catanh subroutine 133
- catanhf subroutine 133
- catanhl subroutine 133
- catanl subroutine 132
- catclose subroutine 134
- catgets subroutine 135
- catopen subroutine 136
- cbrt subroutine 137
- cbrtf subroutine 137
- cbrtl subroutine 137
- ccos, subroutine 138
- ccosf subroutine 138
- ccosh subroutine 139
- ccoshf subroutine 139
- ccoshl subroutine 139
- ccosl subroutine 138

CCSIDs  
  converting 139  
ccsidtocs subroutine 139  
ceil subroutine 140  
ceild128 subroutine 140  
ceild32 subroutine 140  
ceild64 subroutine 140  
ceilf subroutine 140  
ceiling value function  
  ceilf 140  
  ceill 140  
ceill subroutine 140  
cexp subroutine 141  
cexpf subroutine 141  
cexpl subroutine 141  
cfgetospeed subroutine 142  
chacl subroutine 144  
character conversion  
  8-bit processing codes and 575  
  code set converters 531, 532  
  conv subroutines 183  
  Japanese 576  
  Kanji-specific 575  
  multibyte to wide 804, 805  
  translation operations 183  
character manipulation subroutines  
  \_atojis 576  
  \_jistoa 576  
  \_tolower 576  
  \_toupper 576  
  \_tolower 183  
  \_toupper 183  
  atojis 576  
  conv 183  
  ctype 578  
  fgetc 347  
  fputc 1321  
  getc 347  
  getchar 347  
  getw 347  
  isalnum 208  
  isalpha 208  
  isascii 208  
  iscntrl 208  
  isdigit 208  
  isgraph 208  
  isjalnum 578  
  isjalpha 578  
  isjdigit 578  
  isjgraph 578  
  isjhira 578  
  isjis 578  
  isjkanji 578  
  isjkata 578  
  isjlbytekana 578  
  isjlower 578  
  isjparen 578  
  isjprint 578  
  isjpunct 578  
  isjspace 578  
  isjupper 578  
character manipulation subroutines (*continued*)  
  isjxdigit 578  
  islower 208  
  isparent 578  
  isprint 208  
  ispunct 208  
  isspace 208  
  isupper 208  
  isxdigit 208  
  jstoa 576  
  kutentojis 576  
  NCesc 183  
  NCflatchr 183  
  NCtlower 183  
  NCtoNLchar 183  
  NCtoupper 183  
  NCunesc 183  
  putc 1321  
  putchar 1321  
  putw 1321  
  toascii 183  
  tojhira 576  
  tojkata 576  
  tojlower 576  
  toupper 576  
  tolower 183  
  toujis 576  
  toupper 183  
character shaping 684  
character testing  
  isblank 564  
characters  
  classifying 208, 578  
  returning from input streams 347  
  writing to streams 1321  
charsetID  
  multibyte character 193  
chdir subroutine 147  
chmod subroutine 148  
chown subroutine 151  
chownx subroutine 151  
chpass subroutine 154  
chpassx subroutine 156  
chprojattr subroutine 158  
chprojattrdb subroutine 159  
chroot subroutine 160  
chsys subroutine 162  
cimag subroutine 163  
cimagf subroutine 163  
cimagl subroutine 163  
cjstosj subroutine 575  
ckuseracct subroutine 164  
ckuserID subroutine 166  
class subroutine 167  
clearerr macro 270  
clock subroutine 169  
clock subroutines  
  clock\_getcpuclid 169  
  pthread\_condattr\_getclock 1240  
  pthread\_condattr\_setclock 1240  
clock\_getcpuclid subroutine 169

- clock\_getres subroutine 170
- clock\_gettime subroutine 170
- clock\_nanosleep subroutine 172
- clock\_settime subroutine 170
- clog subroutine 174
- clogf subroutine 174
- clogl subroutine 174
- close subroutine 175
- closedir subroutine 943
- closedir64 subroutine 943
- code sets
  - closing converters 531
  - converting names 139
  - opening converters 532
- coded character set IDs
  - converting 139
- command-line flags
  - returning 397
- Common Host Bus Adapter library
  - HBA\_SetRNIDMgmtInfo 523
- compare\_and\_swap subroutine
  - atomic access 176
- compile subroutine 177
- complementary error subroutines
  - erfcl 228
- complex arc cosine subroutines
  - cacos 129
  - cacosf 129
  - cacosl 129
- complex arc hyperbolic cosine subroutines
  - cacosh 130
  - cacoshf 130
  - cacoshl 130
- complex arc hyperbolic sine subroutines
  - casin 132
  - casinf 132
  - casinl 132
- complex arc hyperbolic tangent subroutines
  - catanh 133
  - catanhf 133
  - catanhl 133
- complex arc sine subroutines
  - casin 131
  - casinf 131
  - casinl 131
- complex argument subroutines
  - carg 131
  - cargf 131
  - cargl 131
- complex conjugate subroutines
  - conj 182
  - conjf 182
  - conjl 182
- complex cosine functions
  - ccos 138
  - ccosf 138
  - ccosl 138
- complex exponential functions
  - cexp 141
  - cexpf 141
  - cexpl 141
- complex hyperbolic cosine functions
  - ccosh 139
  - ccoshf 139
  - ccoshl 139
- complex hyperbolic sine subroutines
  - csinh 194
  - csinhf 194
  - csinhl 194
- complex hyperbolic tangent subroutines
  - ctanh 198
  - ctanhf 198
  - ctanhl 198
- complex imaginary functions
  - cimag 163
  - cimagf 163
  - cimagl 163
- complex natural logarithm functions
  - clog 174
  - clogf 174
  - clogl 174
- complex power subroutines
  - cpow 189
  - cpowf 189
  - cpowl 189
- complex projection subroutines
  - cproj 190
  - cprojf 190
  - cprojl 190
- complex tangent functions
  - catan 132
  - catanf 132
  - catanl 132
- Complex tangent subroutines
  - ctan 198
  - ctanf 198
  - ctanl 198
- Computes the base 2 exponential.
  - exp2 247
  - exp2f 247
  - exp2l 247
- confstr subroutine 181
- conj subroutine 182
- conjf subroutine 182
- conjl subroutine 182
- controlling terminal 199
- conv subroutines 183
- conversion
  - date and time representations 206
  - date and time to string representation
    - using asctime subroutine 206
    - using ctime subroutine 206
    - using gmtime subroutine 206
    - using localtime subroutine 206
- converter subroutines
  - btowc 124
  - fwscanf 331
  - iconv\_close 531
  - iconv\_open 532
  - jcode 575
  - mbrlen 792
  - mbrtowc 793

converter subroutines (*continued*)

- mbsinit 797
- mbsrtowcs 802
- swscanf 331
- wscanf 331
- copysign subroutine 185
- copysignd128 subroutine 185
- copysignd32 subroutine 185
- copysignd64 subroutine 185
- copysignf subroutine 185
- copysignl subroutine 185
- core files
  - coredump subroutine 337
  - gencore subroutine 337
- coredump subroutine 337
- cos subroutine 187
- cosf subroutine 187
- cosh subroutine 188
- coshf subroutine 188
- coshl subroutine 188
- cosine subroutines
  - acosf 29
  - acosl 29
  - cosf 187
  - cosl 187
- cosl subroutine 187
- cpow subroutine 189
- cpowf subroutine 189
- cpowl subroutine 189
- cproj subroutine 190
- cprojf subroutine 190
- cprojl subroutine 190
- creal subroutine 190
- crealf subroutine 190
- creall subroutine 190
- creat subroutine 935
- Critical Section Subroutines
  - BeginCriticalSection Subroutine 226
  - EnableCriticalSections Subroutine 226
  - EndCriticalSection Subroutine 226
- crypt subroutine 191
- csid subroutine 193
- csin subroutine 193
- csinf subroutine 193
- csinh subroutine 194
- csinhf subroutine 194
- csinhl subroutine 194
- csinl subroutine 193
- csjtojis subroutine 575
- csjtouj subroutine 575
- csqrt subroutine 195
- csqrtf subroutine 195
- csqrtl subroutine 195
- ctoccsid subroutine 139
- ct\_gen 195
- ct\_hookx 195
- ct\_trcon 197
- ctan subroutine 198
- ctanf subroutine 198
- ctanh subroutine 198
- ctanhf subroutine 198

- ctanh subroutine 198
- ctanf subroutine 198
- ctermid subroutine 199
- ctime subroutine 200
- ctime\_r subroutine 206
- ctime64 subroutine 202
- ctime64\_r subroutine 204
- ctype subroutines 208
- cube root functions
  - cbrtf 137
  - cbrtl 137
- cujtojis subroutine 575
- cujtosj subroutine 575
- current process credentials
  - reading 403
- current process environment
  - reading 405
- current processes
  - getting user name 210
  - group ID
    - initializing 558
    - returning 383
  - path name of controlling terminal 199
  - user ID
    - returning 452
- current working directory
  - getting path name 358
- cursor positions
  - setting 554
- cuserid subroutine 210

## D

- data arrays
  - encrypting 191
- data locks 1026
- data sorting subroutines
  - bsearch 123
  - ftw 324
  - hcreate 527
  - hdestroy 527
  - hsearch 527
  - insque 559
  - lfind 764
  - lsearch 764
  - remque 559
- data space segments
  - changing allocation 122
- date
  - displaying and setting 445
- date format conversions 200
- defect 219851 1262
- defect 220239 412
- defssys subroutine 211
- delssys subroutine 212
- descriptor tables
  - getting size 363
- difftime subroutine 200
- difftime64 subroutine 202
- directories
  - changing 147

- directories *(continued)*
  - changing root 160
  - creating 835
  - directory stream operations 943
  - generating path names 482
  - getting path name of current directory 358
- directory subroutines
  - chdir 147
  - chroot 160
  - closedir 943
  - closedir64 943
  - getcwd 358
  - getwd 479
  - glob 482
  - globfree 485
  - link 718
  - mkdir 835
  - opendir 943
  - opendir64 943
  - readdir 943
  - readdir64 943
  - rewinddir 943
  - rewinddir64 943
  - seekdir 943
  - seekdir64 943
  - telldir 943
  - telldir64 943
- dirname Subroutine 213
- disclaim subroutine 214
- div subroutine 3
- dlclose subroutine 215
- dLError subroutine 216
- dlopen Subroutine 217
- dlsym Subroutine 219
- double precision numbers
  - frexp 315
- drand48 subroutine 221
- drem subroutine 223
- dup subroutine 255
- dup2 subroutine 255

## E

- ecvt subroutine 225
- EnableCriticalSections Subroutine 226
- encrypt subroutine 191
- encryption
  - performing 191
- EndCriticalSection Subroutine 226
- endsent subroutine 369
- endsent\_r subroutine 435
- endgrent subroutine 371
- endhostent subroutine 911
- endpwent subroutine 422
- endrpcent subroutine 427
- endttyent subroutine 450
- endutent subroutine 474
- endvtsent subroutine 476
- environment variables
  - finding default PATH 181
  - finding values 365

- environment variables *(continued)*
  - setting 1326
- erand48 subroutine 221
- erf subroutine 227
- erfc subroutine 228
- erfcf subroutine 228
- erff subroutine 227
- erlog subroutine 229
- erlog\_close subroutine 231
- erlog\_find Subroutines
  - erlog\_find\_first 232
  - erlog\_find\_next 232
  - erlog\_find\_sequence 232
- erlog\_find\_first Subroutine 232
- erlog\_find\_next Subroutine 232
- erlog\_find\_sequence Subroutine 232
- erlog\_open Subroutine 234
- erlog\_set\_direction Subroutine 235
- erlog\_write Subroutine 235
- errorlogging Subroutines
  - erlog\_close 231
  - erlog\_open 234
  - erlog\_set\_direction 235
  - erlog\_write 235
- error functions
  - computing 227
  - erff 227
- error handling
  - math 789
  - returning information 733
- error logs
  - closing 231
  - finding 232
  - opening 234
  - setting direction 235
  - writing 235
  - writing to 229
- error messages
  - placing into program 92
  - writing 1024
- errorlogging subroutines
  - erlog 229
  - perror 1024
- euclidean distance functions
  - hypotf 528
  - hypotl 528
- Euclidean distance functions
  - computing 528
- exec subroutines 236
- execl subroutine 236
- execle subroutine 236
- execlp subroutine 236
- execr subroutine 236
- execution profiling
  - after initialization 849
  - using default data areas 856
  - using defined data areas 851
- execv subroutine 236
- execve subroutine 236
- execvp subroutine 236
- exit subroutine 243

- exp subroutine 245
- exp2 subroutine 247
- exp2f subroutine 247
- exp2l subroutine 247
- expf subroutine 245
- expm1 subroutine 248
- expm1f subroutine 248
- expm1l subroutine 248
- exponential functions
  - computing 245
- exponential subroutines
  - expf 245
  - expm1f, 248
  - expm1l 248
- extended attribute subroutines
  - getea 363
  - listea 724

## F

- f\_hpmgetcounters subroutine 524
- f\_hpmgettimeandcounters subroutine 524
- f\_hpminit subroutine 524
- f\_hpmstart subroutine 524
- f\_hpmstop subroutine 524
- f\_hpmterminate subroutine 524
- f\_hpmtstart subroutine 524
- f\_hpmtstop subroutine 524
- fabs subroutine 249
- fabsd128 subroutine 249
- fabsd32 subroutine 249
- fabsd64 subroutine 249
- fabsf subroutine 249
- fabsl subroutine 249
- faccessx subroutine 4
- fattach Subroutine 250
- fchacl subroutine 144
- fchdir Subroutine 251
- fchmod subroutine 148
- fchown subroutine 151
- fchownx subroutine 151
- fclear subroutine 252
- fclose subroutine 253
- fcntl subroutine 255
- fcvt subroutine 225
- fdetach Subroutine 261
- fdim subroutine 262
- fdimd128 subroutine 262
- fdimd32 subroutine 262
- fdimd64 subroutine 262
- fdimf subroutine 262
- fdiml subroutine 262
- fdopen subroutine 288
- fe\_dec\_getround 263
- fe\_dec\_getround subroutine
  - fe\_dec\_setround 263
- fe\_dec\_setround 263
- feclearexcept subroutine 265
- fegetenv subroutine 265
- fegetexceptflag subroutine 266
- fegetround subroutine 267
- feholdexcept subroutine 267
- feof macro 270
- feraiseexcept subroutine 270
- ferror macro 270
- fesetenv subroutine 265
- fesetexceptflag subroutine 266
- fesetround subroutine 267
- fetch\_and\_add subroutine
  - atomic access 271
- fetch\_and\_and subroutine
  - atomic access 272
- fetch\_and\_or subroutine
  - atomic access 272
- fetestexcept subroutine 273
- feupdateenv subroutine 274
- ffinfo subroutine 274
- fflush subroutine 253
- ffs subroutine 118
- fgetc subroutine 347
- fgetpos subroutine 318
- fgets subroutine 434
- fgetwc subroutine 477
- fgetws subroutine 480
- FIFO files
  - creating 837
- file access permissions
  - changing 144, 148
- file descriptors
  - checking I/O status 1138
  - closing associated files 175
  - controlling 255
  - establishing connections 935
  - performing control functions 561
- file names
  - constructing unique 838
- file ownership
  - changing 151
- file permissions
  - changing 144, 148
- file pointers
  - moving read-write 765
- file subroutines
  - access 4
  - accessx 4
  - dup 255
  - dup2 255
  - endutent 474
  - faccessx 4
  - fclear 252
  - fcntl 255
  - ffinfo 274
  - finfo 274
  - flock 740
  - flockfile 275
  - fpathconf 979
  - fsync 321
  - fsync\_range 321
  - ftrylockfile 275
  - funlockfile 275
  - getc\_unlocked 349
  - getchar\_unlocked 349

file subroutines (*continued*)

- getenv 365
- gettutent 474
- getutid 474
- getutline 474
- lockf 740
- lockfx 740
- lseek 765
- mkfifo 837
- mknod 837
- mkstemp 838
- mktemp 838
- nlist 909
- nlist64 909
- pathconf 979
- pclose 1002
- pipe 1025
- popen 1144
- putc\_unlocked 349
- putchar\_unlocked 349
- putenv 1326
- pututline 474
- setutent 474
- utmpname 474

file system subroutines

- confstr 181
- endfsent 369
- endvfsent 476
- fscntl 316
- getfsent 369
- getfsfile 369
- getfsspec 369
- getfstype 369
- getvfsbyflag 476
- getvfsbyname 476
- getvfsbytype 476
- getvfsent 476
- mntctl 847
- setfsent 369
- setvfsent 476

file systems

- controlling operations 316
- retrieving information 369
- returning mount status 847

file trees

- searching recursively 324

file-implementation characteristics 979

fileno macro 270

files

- binary 310
- closing 175
- creating 837
- creating links 718
- creating space at pointer 252
- determining accessibility 4
- establishing connections 935
- generating path names 482
- getting name list 909
- locking and unlocking 740
- opening 935
- opening streams 288

files (*continued*)

- reading 310
- reading asynchronously 50
- repositioning pointers 318
- revoking access 313
- systems
  - getting information about 435
  - writing asynchronously 61
  - writing binary 310

finfo subroutine 274

finite subroutine 167

finite testing

- isfinite 566

first-in-first-out files 837

flags

- returning 397

floating point multiply-add

- fma 279
- fmaf 279
- fmal 279

floating point numbers

- ldexpf 698, 699
- ldexpl 698, 699
- nextafterf 898
- nextafterl 898
- nexttoward 898
- nexttowardf 898
- nexttowardl 898

floating-point absolute value functions

- computing 277

floating-point environment

- fehldexcept 267
- feupdateenv 274

floating-point environment variables

- fegetenv, 265
- fesetenv 265

floating-point exception

- feraiseexcept 270
- fetestexcept 273

floating-point exceptions 297, 301, 306

- changing floating point status and control register 303

- feclearexcept 265

- flags 295

- querying process state 306

- testing for occurrences 299, 300

floating-point number subroutines

- fdim 262
- fdimf 262
- fdiml 262

floating-point numbers

- converting to strings 225
- determining classifications 167
- fmax 280
- fmaxf 280
- fmaxl 280
- fminf 281
- fminl 281
- fmodf 282
- manipulating 848
- modff 848

floating-point numbers *(continued)*  
  reading and setting rounding modes 302  
  rounding 277

floating-point rounding subroutines  
  nearbyint 897  
  nearbyintf 897  
  nearbyintl 897

floating-point status flags  
  fegetexceptflag 266  
  fesetexceptflag 266

floating-point subroutines 297, 301, 303, 306, 308  
  fp\_sh\_info 303  
  fp\_sh\_trap\_info 303

floating-point trap control 294

flock subroutine 740

flockfile subroutine 275

floor functions  
  floorf 277

floor subroutine 277

floorf subroutine 277

floorl subroutine 277

fma subroutine 279

fmaf subroutine 279

fmal subroutine 279

fmax subroutine 280

fmaxd128 subroutine 280

fmaxd32 subroutine 280

fmaxd64 subroutine 280

fmaxf subroutine 280

fmaxl subroutine 280

fmin subroutine 785

fmind128 subroutine 281

fmind32 subroutine 281

fmind64 subroutine 281

fminf subroutine 281

fminl subroutine 281

fmod subroutine 282

fmodf subroutine 282

fmodl subroutine 282

fmout subroutine 785

fmsg Subroutine 283

fnmatch subroutine 286

fopen subroutine 288

fork subroutine 291

formatted output  
  printing 1168

fp\_any\_enable subroutine 294

fp\_any\_xcp subroutine 299

fp\_clr\_flag subroutine 295

fp\_cpusync subroutine 297

fp\_disable subroutine 294

fp\_disable\_all subroutine 294

fp\_divbyzero subroutine 299

fp\_enable subroutine 294

fp\_enable\_all subroutine 294

fp\_flush\_impresise Subroutine 298

fp\_inexact subroutine 299

fp\_invalid\_op subroutine 299

fp\_iop\_convert subroutine 300

fp\_iop\_infdinf subroutine 300

fp\_iop\_infmzr subroutine 300

fp\_iop\_infsinf subroutine 300

fp\_iop\_invcmp subroutine 300

fp\_iop\_snan subroutine 300

fp\_iop\_sqrt subroutine 300

fp\_iop\_vxsoft subroutine 300

fp\_iop\_zrdzr subroutine 300

fp\_is\_enabled subroutine 294

fp\_overflow subroutine 299

fp\_raise\_xcp subroutine 301

fp\_read\_flag subroutine 295

fp\_read\_rnd subroutine 302

fp\_set\_flag subroutine 295

fp\_sh\_info subroutine 303

fp\_sh\_set\_stat subroutine 303

fp\_sh\_trap\_info subroutine 303

fp\_swap\_flag subroutine 295

fp\_swap\_rnd subroutine 302

fp\_trap subroutine 306

fp\_trapstate subroutine 308

fp\_underflow subroutine 299

fpathconf subroutine 979

fpclassify macro 309

fprintf subroutine 1168

fputc subroutine 1321

fputs subroutine 1331

fputcw subroutine 1339

fputws subroutine 1341

fread subroutine 310

free subroutine 778

free\_agg\_list subroutine 36

freelmb Subroutine 313

fretranlist subroutine 126

freopen subroutine 288

frevoke subroutine 313

frexp subroutine 315

frexpd128 subroutine 314

frexpd32  
  frexpd64 314

frexpd32 subroutine 314

frexpd64 subroutine 314

frexpf subroutine 315

frexpl subroutine 315

fscntl subroutine 316

fseek subroutine 318

fsetpos subroutine 318

fsync subroutine 321

fsync\_range subroutine 321

ftell subroutine 318

ftime subroutine 445

ftok subroutine 322

ftrylockfile subroutine 275

ftw subroutine 324

funlockfile subroutine 275

fwide subroutine 326

fwprintf subroutine 327

fwrite subroutine 310

fwscanf subroutine 331

## G

gai\_strerror subroutine 336



gamma functions  
   computing natural logarithms 336  
 gamma subroutine 336  
 gcd subroutine 785  
 gcvt subroutine 225  
 gencore subroutine 337  
 genpagvalue Subroutine 339  
 get\_malloc\_log subroutine 341  
 get\_malloc\_log\_live subroutine 341  
 get\_speed subroutine 342  
 getargs Subroutine 343  
 getarmlist subroutine 414  
 getaudithostattr, IDtohost, hosttoID, nexthost or  
   putaudithostattr subroutine 345  
 getauthdb subroutine 346  
 getauthdb\_r subroutine 346  
 getc subroutine 347  
 getc\_unlocked subroutine 349  
 getchar subroutine 347  
 getchar\_unlocked subroutine 349  
 getconfattr subroutine 350  
 getconfattrr subroutine 355  
 getcontext or setcontext Subroutine 357  
 getcwd subroutine 358  
 getdate Subroutine 360  
 getdtablesize subroutine 363  
 getea subroutine 363  
 getegid subroutine 370  
 getenv subroutine 365  
 geteuid subroutine 452  
 getevars Subroutine 365  
 getfilehdr subroutine 367  
 getfirstprojdb subroutine 368  
 getfsent subroutine 369  
 getfsent\_r subroutine 435  
 getfsfile subroutine 369  
 getfsspec subroutine 369  
 getfsspec\_r subroutine 435  
 getfstype subroutine 369  
 getfstype\_r subroutine 435  
 getgid subroutine 370  
 getgidx subroutine 370  
 getgrent subroutine 371  
 getgrgid subroutine 371  
 getgrgid\_r subroutine 373  
 getgrnam subroutine 371  
 getgrnam\_r subroutine 374  
 getgroupattr subroutine 375  
 getgroupattrr subroutine 378  
 getgroups subroutine 383  
 getgrpaclattr Subroutine 384  
 gethostent subroutine 910  
 getinterval subroutine 386  
 getiopri 389  
 getitimer subroutine 386  
 getlogin subroutine 393  
 getlogin\_r subroutine 395  
 getlparlist subroutine 414  
 getnextprojdb subroutine 396  
 getopt subroutine 397  
 getpagesize subroutine 399  
 getpaginfo subroutine 400  
 getpagvalue subroutine 401  
 getpagvalue64 subroutine 401  
 getpass subroutine 402  
 getpcred subroutine 403  
 getpeereid subroutine 405  
 getpenv subroutine 405  
 getpgid Subroutine 407  
 getpgrp subroutine 407  
 getpid subroutine 407  
 getportattr Subroutine 408  
 getppid subroutine 407  
 getpri subroutine 411  
 getpriority subroutine 412  
 getproclist subroutine 414  
 getproj subroutine 418  
 getprojdb subroutine 419  
 getprojs subroutine 420  
 getpw Subroutine 421  
 getpwent subroutine 422  
 getpwnam subroutine 422  
 getpwuid subroutine 422  
 getrlimit subroutine 424  
 getrlimit64 subroutine 424  
 getroleattr Subroutine 431  
 getrpcbyname subroutine 427  
 getrpcbynumber subroutine 427  
 getrpcnt subroutine 427  
 getrusage subroutine 428  
 getrusage64 subroutine 428  
 gets subroutine 434  
 getsfile\_r subroutine 435  
 getsid Subroutine 436  
 getssys subroutine 437  
 getsubopt Subroutine 438  
 getsubsvr subroutine 439  
 gettcattr subroutine 440  
 gettimeofday subroutine 445  
 gettimer subroutine 446  
 gettimerid subroutine 449  
 getttyent subroutine 450  
 getttynam subroutine 450  
 getuid subroutine 452  
 getuidx subroutine 452  
 getuinfo subroutine 453  
 getuinfox Subroutine 453  
 getuserattr subroutine 454  
 getuserattrr subroutine 460  
 GetUserAuths Subroutine 467  
 getuserpw subroutine 468  
 getuserpwX subroutine 470  
 getusraclattr Subroutine 472  
 getutent subroutine 474  
 getutid subroutine 474  
 getutline subroutine 474  
 getvfsbyflag subroutine 476  
 getvfsbyname subroutine 476  
 getvfsbytype subroutine 476  
 getvfsent subroutine 476  
 getw subroutine 347  
 getwc subroutine 477

getwchar subroutine 477  
getwd subroutine 479  
getws subroutine 480  
glob subroutine 482  
globfree subroutine 485  
gmtime subroutine 200  
gmtime\_r subroutine 206  
gmtime64 subroutine 202  
gmtime64\_r subroutine 204  
grantpt subroutine 485

## H

hash tables

manipulating 527

HBA subroutines

HBA\_GetEventBuffer 491  
HBA\_GetFC4Statistics 492  
HBA\_GetFCPStatistics 494  
HBA\_GetFcpTargetMappingV2 495  
HBA\_GetPersistentBindingV2 498  
HBA\_OpenAdapterByWWN 503  
HBA\_ScsiInquiryV2 505  
HBA\_ScsiReadCapacityV2 507  
HBA\_ScsiReportLunsV2 509  
HBA\_SendCTPassThruV2 511  
HBA\_SendRLS 515  
HBA\_SendRNIDV2 517  
HBA\_SendRPL 519  
HBA\_SendRPS 520  
HBA\_CloseAdapter Subroutine 486  
HBA\_FreeLibrary Subroutine 487  
HBA\_GetAdapterAttributes Subroutine 487  
HBA\_GetAdapterName Subroutine 490  
HBA\_GetDiscoveredPortAttributes Subroutine 487  
HBA\_GetEventBuffer subroutine 491  
HBA\_GetFC4Statistics subroutine 492  
HBA\_GetFCPStatistics subroutine 494  
HBA\_GetFcpTargetMapping Subroutine 496  
HBA\_GetFcpTargetMappingV2 subroutine 495  
HBA\_GetNumberOfAdapters Subroutine 497  
HBA\_GetPersistentBinding Subroutine 493  
HBA\_GetPersistentBindingV2 subroutine 498  
HBA\_GetPortAttributes Subroutine 487  
HBA\_GetPortAttributesByWWN Subroutine 487  
HBA\_GetPortStatistics Subroutine 499  
HBA\_GetRNIDMgmtInfo Subroutine 500  
HBA\_GetVersion Subroutine 501  
HBA\_LoadLibrary Subroutine 502  
HBA\_OpenAdapter Subroutine 503  
HBA\_OpenAdapterByWWN subroutine 503  
HBA\_RefreshInformation Subroutine 504  
HBA\_ScsiInquiryV2 subroutine 505  
HBA\_ScsiReadCapacityV2 subroutine 507  
HBA\_ScsiReportLunsV2 subroutine 509  
HBA\_SendCTPassThru Subroutine 510  
HBA\_SendCTPassThruV2 subroutine 511  
HBA\_SendReadCapacity Subroutine 512  
HBA\_SendReportLUNs Subroutine 513  
HBA\_SendRLS subroutine 515  
HBA\_SendRNID Subroutine 516

HBA\_SendRNIDV2 subroutine 517  
HBA\_SendRPL subroutine 519  
HBA\_SendRPS subroutine 520  
HBA\_SendScsiInquiry Subroutine 521  
HBA\_SetRNIDMgmtInfo Subroutine 523  
hcreate subroutine 527  
hdestroy subroutine 527  
Host Bus Adapter API  
HBA\_CloseAdapter 486  
HBA\_FreeLibrary 487  
HBA\_GetAdapterAttributes 487  
HBA\_GetAdapterName 490  
HBA\_GetDiscoveredPortAttributes 487  
HBA\_GetFcpPersistentBinding 493  
HBA\_GetFcpTargetMapping 496  
HBA\_GetNumberOfAdapters 497  
HBA\_GetPortAttributes 487  
HBA\_GetPortAttributesByWWN 487  
HBA\_GetPortStatistics 499  
HBA\_GetRNIDMgmtInfo 500  
HBA\_GetVersion 501  
HBA\_LoadLibrary 502  
HBA\_OpenAdapter 503  
HBA\_RefreshInformation 504  
HBA\_SendCTPassThru 510  
HBA\_SendReadCapacity 512  
HBA\_SendReportLUNs 513  
HBA\_SendRNID 516  
HBA\_SendScsiInquiry 521  
HBA\_SetRNIDMgmtInfo 523  
hpmGetCounters subroutine 524  
hpmGetTimeAndCounters subroutine 524  
hpmInit subroutine 524  
hpmStart subroutine 524  
hpmStop subroutine 524  
hpmTerminate subroutine 524  
hpmTstart subroutine 524  
hpmTstop subroutine 524  
hsearch subroutine 527  
hyperbolic cosine subroutines  
coshf 188  
coshl 188  
hypot subroutine 528  
hypotf subroutine 528  
hypotl subroutine 528

## I

I/O asynchronous subroutines

aio\_cancel 38  
aio\_error 42  
aio\_fsync 45  
aio\_nwait 47  
aio\_nwait\_timeout 49  
aio\_read 50  
aio\_return 55  
aio\_suspend 58  
aio\_write 61  
lio\_listio 720  
poll 1138

I/O low-level subroutines 175, 935

I/O low-level subroutines (*continued*)

- creat 935
- open 935
- I/O requests
  - canceling 38
  - listing 720
  - retrieving error status 42
  - retrieving return status 55
  - suspending 58
- I/O stream macros
  - clearerr 270
  - feof 270
  - ferror 270
  - fileno 270
- I/O stream subroutines
  - fclose 253
  - fdopen 288
  - fflush 253
  - fgetc 347
  - fgetpos 318
  - fgets 434
  - fgetwc 477
  - fgetws 480
  - fopen 288
  - fprintf 1168
  - fputc 1321
  - fputs 1331
  - fputwc 1339
  - fputws 1341
  - fread 310
  - freopen 288
  - fseek 318
  - fsetpos 318
  - ftell 318
  - fwide 326
  - fwprintf 327
  - fwrite 310
  - getc 347
  - getchar 347
  - gets 434
  - getw 347
  - getwc 477
  - getwchar 477
  - getws 480
  - printf 1168
  - putc 1321
  - putchar 1321
  - puts 1331
  - putw 1321
  - putwc 1339
  - putwchar 1339
  - putws 1341
  - rewind 318
  - sprintf 1168
  - swprintf 327
  - vfprintf 1168
  - vprintf 1168
  - vsprintf 1168
  - vwsprintf 1168
  - wprintf 327
  - wsprintf 1168

I/O terminal subroutines

- cfsetispeed 142
- ioctl 561
- ioctl32 561
- ioctl32x 561
- ioctlx 561
- iconv\_close subroutine 531
- iconv\_open subroutine 532
- identification subroutines
  - endgrent 371
  - endpwent 422
  - getconfattr 350
  - getgrent 371
  - getgrgid 371
  - getgrnam 371
  - getgroupattr 375
  - getpwent 422
  - getpwnam 422
  - getpwuid 422
  - gettcattr 440
  - getuinfo 453
  - getuserattr 350, 454
  - IDtgroup 375
  - IDtouser 454
  - nextgroup 375
  - nextuser 454
  - putconfattr 350
  - putgroupattr 375
  - putpwent 422
  - puttcattr 440
  - putuserattr 454
  - setgrent 371
  - setpwent 422
- idpthreadsa 191
- IDtgroup subroutine 375
- IDtouser subroutine 454
- IEE Remainders
  - computing 223
- ilogb subroutine 534
- ilogbd128 subroutine 747
- ilogbd32 subroutine 747
- ilogbd64 subroutine 747
- ilogbf subroutine 534
- ilogbl subroutine 534
- IMAIXMapping subroutine 536
- IMAuxCreate callback subroutine 537
- IMAuxDestroy callback subroutine 537
- IMAuxDraw callback subroutine 538
- IMAuxHide callback subroutine 539
- imaxabs subroutine 535
- imaxdiv subroutine 535
- IMBeep callback subroutine 539
- IMClose subroutine 540
- IMCreate subroutine 541
- IMDestroy subroutine 541
- IMFilter subroutine 542
- IMFreeKeymap subroutine 543
- IMIndicatorDraw callback subroutine 543
- IMIndicatorHide callback subroutine 544
- IMInitialize subroutine 545
- IMInitializeKeymap subroutine 546

- IMIoctl subroutine 547
- IMLookupString subroutine 549
- IMProcess subroutine 550
- IMProcessAuxiliary subroutine 551
- IMQueryLanguage subroutine 552
- IMSimpleMapping subroutine 553
- IMTextCursor callback subroutine 554
- IMTextDraw callback subroutine 555
- IMTextHide callback subroutine 555
- IMTextStart callback subroutine 556
- imul\_dbl subroutine 3
- incinterval subroutine 386
- inet\_aton subroutine 557
- infinity values
  - isinf 568
- initgroups subroutine 558
- initialize subroutine 558
- input method
  - checking language support 552
  - closing 540
  - control and query operations 547
  - creating instance 541
  - destroying instance 541
  - initializing for particular language 545
- input method keymap
  - initializing 543, 546
  - mapping key and state pair to string 536, 549, 553
- input method subroutines
  - callback functions
    - IMAuxCreate 537
    - IMAuxDestroy 537
    - IMAuxDraw 538
    - IMAuxHide 539
    - IMBeep 539
    - IMIndicatorDraw 543
    - IMIndicatorHide 544
    - IMTextCursor 554
    - IMTextDraw 555
    - IMTextHide 555
    - IMTextStart 556
  - IMAIXMapping 536
  - IMClose 540
  - IMCreate 541
  - IMDestroy 541
  - IMFilter 542
  - IMFreeKeymap 543
  - IMinitialize 545
  - IMInitializeKeymap 546
  - IMIoctl 547
  - IMLookupString 549
  - IMProcess 550
  - IMProcessAuxiliary 551
  - IMQueryLanguage 552
  - IMSimpleMapping 553
- input streams
  - reading character string from 480
  - reading single character from 477
  - returning characters or words 347
- insque subroutine 559
- install\_lwcf\_handler() subroutine 560
- integers
  - computing absolute values 3
  - computing division 3
  - computing double-precision multiplication 3
  - performing arithmetic 785
- Internet addresses
  - converting to ASCII strings 557
- interoperability subroutines
  - ccsidtocs 139
  - ctoccsid 139
- interprocess channels
  - creating 1025
- interprocess communication keys 322
- interval timers
  - allocating per process 449
  - manipulating expiration time 386
  - returning values 386
- inverse hyperbolic cosine subroutines
  - acoshf 30
  - acoshl 30
- inverse hyperbolic functions
  - computing 30, 90
- inverse hyperbolic sine subroutines
  - asinhf 90
  - asinhl 90
- inverse hyperbolic tangent subroutines
  - atanhf 95
  - atanhl 95
- invert subroutine 785
- ioctl subroutine 561
- ioctl32 subroutine 561
- ioctl32x subroutine 561
- ioctlx subroutine 561
- is\_wctype subroutine 574
- isalnum subroutine 208
- isalpha subroutine 208
- isascii subroutine 208
- isblank subroutine 564
- iscntrl subroutine 208
- isdigit subroutine 208
- isendwin Subroutine 565
- isfinite macro 566
- isgraph subroutine 208
- isgreater macro 566
- isgreaterequal subroutine 567
- isinf subroutine 568
- isless macro 568
- islessequal macro 569
- islessgreater macro 570
- islower subroutine 208
- isnan subroutine 167
- isnormal macro 570
- isprint subroutine 208
- ispunct subroutine 208
- isspace subroutine 208
- isunordered macro 571
- isupper subroutine 208
- iswalnum subroutine 572
- iswalpha subroutine 572
- iswblank subroutine 573
- iswcntrl subroutine 572

iswctype subroutine 574  
iswdigit subroutine 572  
iswgraph subroutine 572  
iswlower subroutine 572  
iswprint subroutine 572  
iswpunct subroutine 572  
iswspace subroutine 572  
iswupper subroutine 572  
iswxdigit subroutine 572  
isxdigit subroutine 208  
itom subroutine 785  
itrunc subroutine 277

## J

j0 subroutine 119  
j1 subroutine 119  
Japanese conv subroutines 576  
Japanese ctype subroutines 578  
jcode subroutines 575  
JFS  
    controlling operations 316  
JIS character conversions 575  
jistoa subroutine 576  
jistosj subroutine 575  
jistouj subroutine 575  
jn subroutine 119  
Journaled File System 255  
jrand48 subroutine 221

## K

Kanji character conversions 575  
keyboard events  
    processing 542, 550  
kill subroutine 580  
killpg subroutine 580  
kleanup subroutine 581  
knlist subroutine 582  
kpidstate subroutine 584  
kutentojis subroutine 576

## L

l3tol subroutine 586  
l64a subroutine 1  
l64a\_r subroutine 587  
labs subroutine 3  
LAPI\_Addr\_get subroutine 588  
LAPI\_Addr\_set subroutine 589  
LAPI\_Address subroutine 591  
LAPI\_Address\_init subroutine 592  
LAPI\_Address\_init64 594  
LAPI\_Amsend subroutine 596  
LAPI\_Amsendv subroutine 601  
LAPI\_Fence subroutine 609  
LAPI\_Get subroutine 610  
LAPI\_Getcptr subroutine 613  
LAPI\_Getv subroutine 614  
LAPI\_Gfence subroutine 618  
LAPI\_Init subroutine 619

LAPI\_Msg\_string subroutine 624  
LAPI\_Msgpoll subroutine 626  
LAPI\_Nopoll\_wait subroutine 628  
LAPI\_Probe subroutine 629  
LAPI\_Purge\_totask subroutine 630  
LAPI\_Put subroutine 632  
LAPI\_Putv subroutine 634  
LAPI\_Qenv subroutine 638  
LAPI\_Resume\_totask subroutine 641  
LAPI\_Rmw subroutine 642  
LAPI\_Rmw64 subroutine 646  
LAPI\_Senv subroutine 650  
LAPI\_Setcptr subroutine 652  
LAPI\_Setcptr\_wstatus subroutine 654  
LAPI\_Term subroutine 655  
LAPI\_Util subroutine 657  
LAPI\_Waitcptr subroutine 668  
LAPI\_Xfer structure types 671  
LAPI\_Xfer subroutine 670  
lapi\_xfer\_type\_t 671  
layout values  
    querying 687  
    setting 689  
    transforming text 692  
LayoutObject  
    creating 683  
    freeing 695  
lcong48 subroutine 221  
ldaclose subroutine 697  
ldahread subroutine 696  
ldaopen subroutine 707  
ldclose subroutine 697  
ldexp subroutine 699  
ldexpd128 subroutine 698  
ldexpd32 subroutine 698  
ldexpd64 subroutine 698  
ldexpf subroutine 699  
ldexpl subroutine 699  
ldfhread subroutine 701  
ldgetname subroutine 702  
ldiv subroutine 3  
ldlinit subroutine 704  
ldlitem subroutine 704  
ldlnseek subroutine 705  
ldlread subroutine 704  
ldlseek subroutine 705  
ldnrseek subroutine 709  
ldnshread subroutine 710  
ldnsseek subroutine 712  
ldohseek subroutine 707  
ldopen subroutine 707  
ldrseek subroutine 709  
ldshread subroutine 710  
ldsseek subroutine 712  
ldtbindex subroutine 713  
ldtbread subroutine 714  
ldtbseek subroutine 715  
lfind subroutine 764  
lgamma subroutine 716  
lgammaf subroutine 716  
lgammal subroutine 716

libhpm subroutines  
  f\_hpmgetcounters 524  
  f\_hpmgettimeandcounters 524  
  f\_hpminit 524  
  f\_hpmstart 524  
  f\_hpmstop 524  
  f\_hpmterminate 524  
  f\_hpmtstart 524  
  f\_hpmtstop 524  
  hpmGetCounters 524  
  hpmGetTimeAndCounters 524  
  hpmInit 524  
  hpmStart 524  
  hpmStop 524  
  hpmTerminate 524  
  hpmTstart 524  
  hpmTstop 524  
linear searches 764  
lineout subroutine 717  
link subroutine 718  
lio\_listio subroutine 720  
listea subroutine 724  
llabs subroutine 3  
lldiv subroutine 3  
llrint subroutine 725  
llrintd128 subroutine 725  
llrintd32 subroutine 725  
llrintd64 subroutine 725  
llrintf subroutine 725  
llrintl subroutine 725  
llround subroutine 726  
llroundd128 subroutine 726  
llroundd32 subroutine 726  
llroundd64 subroutine 726  
llroundf subroutine 726  
llroundl subroutine 726  
load subroutine 727  
loadAndInit 727  
loadbind subroutine 731  
loadquery subroutine 733  
locale subroutines  
  localeconv 735  
  nl\_langinfo 907  
locale-dependent conventions 735  
localeconv subroutine 735  
locales  
  returning language information 907  
localtime subroutine 200  
localtime\_r subroutine 206  
localtime64 subroutine 202  
localtime64\_r subroutine 204  
lockf subroutine 740  
lockfx subroutine 740  
log gamma functions  
  lgamma 716  
  lgammaf 716  
  lgammal 716  
log subroutine 749  
log10 subroutine 743  
log10f subroutine 743  
log1p subroutine 744  
log1pf subroutine 744  
log1pl subroutine 744  
log2 subroutine 745  
log2f subroutine 745  
log2l subroutine 745  
logarithmic functions  
  computing 245  
logb subroutine 748  
logbd128 subroutine 746  
logbd32 subroutine 746  
logbd64 subroutine 746  
logbf subroutine 748  
logbl subroutine 748  
logf subroutine 749  
logical volumes  
  querying 766  
login name  
  getting 393, 395  
loginfailed Subroutine 750  
loginrestrictions Subroutine 752  
loginrestrictionsx subroutine 754  
loginsuccess Subroutine 757  
long integers  
  converting to strings 587  
long integers, converting  
  to 3-byte integers 586  
  to base-64 ASCII strings 1  
lpar\_get\_info subroutine 759  
lpar\_set\_resources subroutine 760  
lrand48 subroutine 221  
lrint subroutine 761  
lrintd128 subroutine 761  
lrintd32 subroutine 761  
lrintd64 subroutine 761  
lrintf subroutine 761  
lrintl subroutine 761  
lround subroutine 763  
lroundd128 subroutine 763  
lroundd32 subroutine 763  
lroundd64 subroutine 763  
lroundf subroutine 763  
lroundl subroutine 763  
lsearch subroutine 764  
lseek subroutine 765  
lto3 subroutine 586  
LVM logical volume subroutines  
  lvm\_querylv 766  
LVM physical volume subroutines  
  lvm\_querypv 770  
LVM volume group subroutines  
  lvm\_queryvg 774  
  lvm\_queryvgs 777  
lvm\_querylv subroutine 766  
lvm\_querypv subroutine 770  
lvm\_queryvg subroutine 774  
lvm\_queryvgs subroutine 777

## M

m\_in subroutine 785  
m\_out subroutine 785

- macros
  - assert 92
- madd subroutine 785
- madvise subroutine 787
- makecontext Subroutine 788
- mallinfo subroutine 778
- mallinfo\_heap subroutine 778
- malloc subroutine 778
- mallopt subroutine 778
- mapped files
  - synchronizing 890
- MatchAllAuths Subroutine 790
- MatchAllAuthsList Subroutine 790
- MatchAnyAuthsList Subroutine 790
- math errors
  - handling 789
- matherr subroutine 789
- mblen subroutine 791
- mbrlen subroutine 792
- mbrtowc subroutine 793
- mbsadvance subroutine 794
- mbscat subroutine 795
- mbschr subroutine 796
- mbscmp subroutine 795
- mbscpy subroutine 795
- mbsinit subroutine 797
- mbsinvalid subroutine 798
- mbslen subroutine 798
- mbsncat subroutine 799
- mbsncmp subroutine 799
- mbsncpy subroutine 799
- mbspbrk subroutine 800
- mbsrchr subroutine 801
- mbsrtowcs subroutine 802
- mbstomb subroutine 803
- mbstowcs subroutine 804
- mbswidth subroutine 805
- mbtowc subroutine 805
- ncmp subroutine 785
- mdiv subroutine 785
- memccpy subroutine 807
- memchr subroutine 807
- memcmp subroutine 807
- memcpy subroutine 807
- memmove subroutine 807
- memory allocation 778
- memory area operations 807
- memory management
  - controlling execution profiling 849, 851, 856
  - defining addresses 224
  - defining available paging space 1181
  - disclaiming memory content 214
  - generating IPC keys 322
  - returning system page size 399
- memory management subroutines
  - alloca 778
  - calloc 778
  - disclaim 214
  - free 778
  - ftok 322
  - gai\_strerror 336
- memory management subroutines *(continued)*
  - getpagesize 399
  - madvise 787
  - mallinfo 778
  - mallinfo\_heap 778
  - malloc 778
  - mallopt 778
  - memccpy 807
  - memchr 807
  - memcmp 807
  - memcpy 807
  - memmove 807
  - memset 807
  - mincore 808
  - mmap 843
  - moncontrol 849
  - monitor 851
  - monstartup 856
  - mprotect 859
  - msem\_init 875
  - msem\_lock 876
  - msem\_remove 877
  - msem\_unlock 878
  - msleep 889
  - msync 890
  - munmap 893
  - mwakeup 894
  - psdanger 1181
  - realloc 778
- memory mapping
  - advising system of paging behavior 787
  - determining page residency status 808
  - file-system objects 843
  - modifying access protections 859
  - putting a process to sleep 889
  - semaphores
    - initializing 875
    - locking 876
    - removing 877
    - unlocking 878
  - synchronizing mapped files 890
  - unmapping regions 893
  - waking a process 894
- memory pages
  - determining residency 808
- memory semaphores
  - initializing 875
  - locking 876
  - putting a process to sleep 889
  - removing 877
  - unlocking 878
  - waking a process 894
- memory subroutines
  - allocmb 68
  - freelmb 313
- memset subroutine 807
- message catalogs
  - closing 134
  - opening 136
  - retrieving messages 135
- message control operations 879

message facility subroutines  
   catclose 134  
   catgets 135  
   catopen 136  
 message queue identifiers 881  
 message queue subroutines  
   mq\_receive 870  
   mq\_send 872  
   mq\_timedreceive 870  
   mq\_timedsend 872  
 message queues  
   checking I/O status 1138  
   reading messages from 882  
   receiving messages from 887  
   sending messages to 885  
 Micro-Partitioning  
   lpar\_get\_info 759  
 min subroutine 785  
 mincore subroutine 808  
 mkdir subroutine 835  
 mkfifo subroutine 837  
 mknod subroutine 837  
 mkstemp subroutine 838  
 mktemp subroutine 838  
 mktime subroutine 200  
 mktime64 subroutine 202  
 mlockall subroutine 840, 841  
 mmap subroutine 843  
 mntctl subroutine 847  
 modf subroutine 848  
 modff subroutine 848  
 modfl subroutine 848  
 modulo remainders  
   computing 277  
 moncontrol subroutine 849  
 monitor subroutine 851  
 monstartup subroutine 856  
 mout subroutine 785  
 move subroutine 785  
 mprotect subroutine 859  
 mq\_close subroutine 861  
 mq\_getattr subroutine 862  
 mq\_notify subroutine 863  
 mq\_open subroutine 864  
 mq\_receive subroutine 866, 870  
 mq\_send subroutine 868, 872  
 mq\_setattr subroutine 869  
 mq\_timedreceive subroutine 870  
 mq\_timedsend subroutine 872  
 mq\_unlink subroutine 874  
 mrand48 subroutine 221  
 msem\_init subroutine 875  
 msem\_lock subroutine 876  
 msem\_remove subroutine 877  
 msem\_unlock subroutine 878  
 msgctl subroutine 879  
 msgget subroutine 881  
 msgrcv subroutine 882  
 msgsnd subroutine 885  
 msgxrcv subroutine 887  
 msleep subroutine 889  
 msqrt subroutine 785  
 msub subroutine 785  
 msync subroutine 890  
 mt\_trce() subroutine 891  
 mult subroutine 785  
 multibyte character subroutines  
   csid 193  
   mblen 791  
   mbsadvance 794  
   mbscat 795  
   mbschr 796  
   mbscmp 795  
   mbscpy 795  
   mbsinvalid 798  
   mbslen 798  
   mbsncat 799  
   mbsncmp 799  
   mbsncpy 799  
   mbspbrk 800  
   mbsrchr 801  
   mbstomb 803  
   mbstowcs 804  
   mbswidth 805  
   mbtowc 805  
 multibyte characters  
   converting to wide 804, 805  
   determining display width of 805  
   determining length of 791  
   determining number of 798  
   extracting from string 803  
   locating character sequences 800  
   locating next character 794  
   locating single characters 796, 801  
   operations on null-terminated strings 795, 799  
   returning charsetID 193  
   validating 798  
 munlockall subroutine 840, 841  
 munmap subroutine 893  
 mwakeup subroutine 894

## N

### NaN

nan 895  
 nanf 895  
 nanl 895  
 nan subroutine 895  
 nand128 subroutine 895  
 nand32 subroutine 895  
 nand64 subroutine 895  
 nanf subroutine 895  
 nanl subroutine 895  
 nanosleep subroutine 896  
 natural logarithm functions  
   logf 749  
   logl 749  
 natural logarithms  
   log1pf 744  
   log1pl 744  
 NCesc subroutine 183  
 NCflatchr subroutine 183



NCtlower subroutine 183  
 NCtoNLchar subroutine 183  
 NCtoupper subroutine 183  
 NCunesc subroutine 183  
 nearbyint subroutine 897  
 nearbyintd128 subroutine 897  
 nearbyintd32 subroutine 897  
 nearbyintd64 subroutine 897  
 nearbyintf subroutine 897  
 nearbyintl subroutine 897  
 nearest subroutine 277  
 network host entries  
     retrieving 910  
 new-process image file 236  
 newpass subroutine 901  
 newpassx subroutine 903  
 nextafter subroutine 898  
 nextafterd128 Subroutine 900  
 nextafterd32 Subroutine 900  
 nextafterd64 Subroutine 900  
 nextafterf subroutine 898  
 nextafterl subroutine 898  
 nextgroup subroutine 375  
 nextgrpacl Subroutine 384  
 nextrole Subroutine 431  
 nexttoward subroutine 898  
 nexttowardd128 Subroutine 900  
 nexttowardd32 subroutine 900  
 nexttowardd64 Subroutine 900  
 nexttowardf subroutine 898  
 nexttowardl subroutine 898  
 nextuser subroutine 454  
 nextusracl Subroutine 472  
 nftw subroutine 905  
 nice subroutine 412  
 nl\_langinfo subroutine 907  
 nlist subroutine 909  
 nlist64 subroutine 909  
 nrand48 subroutine 221  
 number manipulation function  
     copysignd128 185  
     copysignd32 185  
     copysignd64 185  
     copysignf 185  
     copysignl 185  
 numbers  
     generating  
         pseudo-random 221  
 numerical manipulation subroutines 336  
     a64l 1  
     abs 3  
     acos 29  
     acosf 29  
     acosh 30  
     acosl 29  
     asin 91  
     asinh 90  
     asinl 91  
     atan 94  
     atan2 93  
     atan2f 93

numerical manipulation subroutines (*continued*)  
     atan2l 93  
     atanf 94  
     atanh 95  
     atanhf 95  
     atanhl 95  
     atanl 94  
     atof 96  
     atoff 96  
     atol 98  
     atoll 98  
     cabs 528  
     cbrt 137  
     ceil 140  
     ceild128 140  
     ceild32 140  
     ceild64 140  
     ceilf 140  
     ceill 140  
     class 167  
     cos 187  
     div 3  
     drand48 221  
     drem 223  
     ecvt 225  
     erand48 221  
     erf 227  
     erfc 228  
     exp 245  
     expm1 248  
     fabs 249  
     fabsl 249  
     fcvt 225  
     finite 167  
     flood128 277  
     flood32 277  
     flood64 277  
     floor 277  
     floorl 277  
     fmin 785  
     fmod 282  
     fmodl 282  
     fp\_any\_enable 294  
     fp\_any\_xcp 299  
     fp\_clr\_flag 295  
     fp\_disable 294  
     fp\_disable\_all 294  
     fp\_divbyzero 299  
     fp\_enable 294  
     fp\_enable\_all 294  
     fp\_inexact 299  
     fp\_invalid\_op 299  
     fp\_iop\_convert 300  
     fp\_iop\_infdinf 300  
     fp\_iop\_infmzr 300  
     fp\_iop\_infsinf 300  
     fp\_iop\_invcmp 300  
     fp\_iop\_snan 300  
     fp\_iop\_sqrt 300  
     fp\_iop\_zrdzr 300  
     fp\_is\_enabled 294

numerical manipulation subroutines (*continued*)

fp\_overflow 299  
fp\_read\_flag 295  
fp\_read\_rnd 302  
fp\_set\_flag 295  
fp\_swap\_flag 295  
fp\_swap\_rnd 302  
fp\_underflow 299  
frexp 315  
frexpl 315  
gamma 336  
gcd 785  
gcvl 225  
hypot 528  
ilogb 534  
imul\_dbl 3  
invert 785  
isnan 167  
itom 785  
itrunc 277  
j0 119  
j1 119  
jn 119  
jrand48 221  
l3tol 586  
l64a 1  
labs 3  
lcong48 221  
ldexp 698, 699  
ldexpl 698, 699  
ldiv 3  
llabs 3  
lldiv 3  
log 749  
log10 743  
log1p 744  
logb 746, 748  
lrand48 221  
ltol3 586  
m\_in 785  
m\_out 785  
madd 785  
matherr 789  
mcmp 785  
mdiv 785  
min 785  
modf 848  
modfl 848  
mout 785  
move 785  
mrnd48 221  
msqrt 785  
msub 785  
mult 785  
nearest 277  
nextafter 898  
nrnd48 221  
omin 785  
omout 785  
pow 785, 1166  
rpow 785

numerical manipulation subroutines (*continued*)

sdiv 785  
seed48 221  
srand48 221  
trunc 277  
uitrunc 277  
umul\_dbl 3  
unordered 167  
y0 119  
y1 119  
yn 119

## O

Object Data Manager 923

object file access subroutines

ldaclose 697  
ldahread 696  
ldaopen 707  
ldclose 697  
ldfhread 701  
ldgetname 702  
ldlinit 704  
ldlitem 704  
ldlread 704  
ldlseek 705  
ldnlseek 705  
ldnrseek 709  
ldnshread 710  
ldnsseek 712  
ldohseek 707  
ldopen 707  
ldrseek 709  
ldshread 710  
ldsseek 712  
ldtbindex 713  
ldtbread 714  
ldtbseek 715

object file subroutines

load 727  
loadbind 731  
loadquery 733

object files

closing 697  
computing symbol table entries 713  
controlling run-time resolution 731  
listing 733  
loading and binding 727  
manipulating line number entries 704  
providing access 707  
reading archive headers 696  
reading file headers 701  
reading indexed section headers 710  
reading symbol table entries 714  
retrieving symbol names 702  
seeking to indexed sections 712  
seeking to line number entries 705  
seeking to optional file header 707  
seeking to relocation entries 709  
seeking to symbol tables 715

- objects
  - setting locale-dependent conventions 735
- ODM
  - ending session 933
  - error message strings 916
  - freeing memory 917
- ODM (Object Data Manager)
  - initializing 923
  - running specified method 931
- ODM object classes
  - adding objects 912
  - changing objects 913
  - closing 914
  - creating 915
  - locking 923
  - opening 926
  - removing 928
  - removing objects 927, 929
  - retrieving class symbol structures 925
  - retrieving objects 918, 919, 921
  - setting default path location 932
  - setting default permissions 933
  - unlocking 934
- ODM subroutines
  - odm\_add\_obj 912
  - odm\_change\_obj 913
  - odm\_close\_class 914
  - odm\_create\_class 915
  - odm\_err\_msg 916
  - odm\_free\_list 917
  - odm\_get\_by\_id 918
  - odm\_get\_first 921
  - odm\_get\_list 919
  - odm\_get\_next 921
  - odm\_get\_obj 921
  - odm\_initialize 923
  - odm\_lock 923
  - odm\_mount\_class 925
  - odm\_open\_class 926
  - odm\_open\_class\_ronly 926
  - odm\_rm\_by\_id 927
  - odm\_rm\_class 928
  - odm\_rm\_obj 929
  - odm\_run\_method 931
  - odm\_set\_path 932
  - odm\_set\_perms 933
  - odm\_terminate 933
  - odm\_unlock 934
- odm\_add\_obj subroutine 912
- odm\_change\_obj subroutine 913
- odm\_close\_class subroutine 914
- odm\_create\_class subroutine 915
- odm\_err\_msg subroutine 916
- odm\_free\_list subroutine 917
- odm\_get\_by\_id subroutine 918
- odm\_get\_first subroutine 921
- odm\_get\_list subroutine 919
- odm\_get\_next subroutine 921
- odm\_get\_obj subroutine 921
- odm\_initialize subroutine 923
- odm\_lock subroutine 923

- odm\_mount\_class subroutine 925
- odm\_open\_class subroutine 926
- odm\_open\_class\_ronly subroutine 926
- odm\_rm\_by\_id subroutine 927
- odm\_rm\_class subroutine 928
- odm\_rm\_obj subroutine 929
- odm\_run\_method subroutine 931
- odm\_set\_path subroutine 932
- odm\_set\_perms subroutine 933
- odm\_terminate subroutine 933
- odm\_unlock subroutine 934
- omin subroutine 785
- omout subroutine 785
- open file descriptors
  - controlling 255
  - performing control functions 561
- open subroutine
  - described 935
- opendir subroutine 943
- opendir64 subroutine 943
- openx subroutine
  - described 935
- output stream
  - writing character string to 1341
  - writing single character to 1339

## P

- PAG Services
  - genpagvalue 339
- paging memory
  - behavior 787
  - defining available space 1181
- PAM subroutines
  - pam\_acct\_mgmt 946
  - pam\_authenticate 947
  - pam\_chauthtok 949
  - pam\_close\_session 950
  - pam\_end 951
  - pam\_get\_data 952
  - pam\_get\_item 953
  - pam\_get\_user 954
  - pam\_getenv 955
  - pam\_getenvlist 956
  - pam\_open\_session 957
  - pam\_putenv 958
  - pam\_set\_data 959
  - pam\_set\_item 960
  - pam\_setcred 961
  - pam\_sm\_acct\_mgmt 963
  - pam\_sm\_authenticate 964
  - pam\_sm\_chauthtok 965
  - pam\_sm\_close\_session 967
  - pam\_sm\_open\_session 968
  - pam\_sm\_setcred 969
  - pam\_start 970
  - pam\_strerror 973
- pam\_acct\_mgmt subroutine 946
- pam\_authenticate subroutine 947
- pam\_chauthtok subroutine 949
- pam\_close\_session subroutine 950

pam\_end subroutine 951  
 pam\_get\_data subroutine 952  
 pam\_get\_item subroutine 953  
 pam\_get\_user subroutine 954  
 pam\_getenv subroutine 955  
 pam\_getenvlist subroutine 956  
 pam\_open\_session subroutine 957  
 pam\_putenv subroutine 958  
 pam\_set\_data subroutine 959  
 pam\_set\_item subroutine 960  
 pam\_setcred subroutine 961  
 pam\_sm\_acct\_mgmt subroutine 963  
 pam\_sm\_authenticate subroutine 964  
 pam\_sm\_chauthtok subroutine 965  
 pam\_sm\_close\_session subroutine 967  
 pam\_sm\_open\_session subroutine 968  
 pam\_sm\_setcred subroutine 969  
 pam\_start subroutine 970  
 pam\_strerror subroutine 973  
 passwdexpired 973  
 passwdexpiredx subroutine 974  
 passwdpolicy subroutine 976  
 passwdstrength subroutine 978  
 password maintenance  
     password changing 154  
 password subroutines  
     passwdpolicy 976  
     passwdstrength 978  
 passwords  
     generating new 901  
     reading 402  
 pathconf subroutine 979  
 pause subroutine 982  
 pcap\_close 982  
 pcap\_compile 983  
 pcap\_datalink 983  
 pcap\_dispatch 984  
 pcap\_dump 985  
 pcap\_dump\_close 986  
 pcap\_dump\_open 987  
 pcap\_file 988  
 pcap\_fileno 988  
 pcap\_geterr 989  
 pcap\_is\_swapped 990  
 pcap\_lookupdev 990  
 pcap\_lookupnet 991  
 pcap\_loop 992  
 pcap\_major\_version 993  
 pcap\_next 994  
 pcap\_open\_live 995  
 pcap\_open\_live\_sb  
     pcap\_open\_live 996  
 pcap\_open\_live\_sb Subroutine 996  
 pcap\_open\_offline 996  
 pcap\_perror 998  
 pcap\_setfilter 999  
 pcap\_snapshot 1000  
 pcap\_stats 1000  
 pcap\_strerror 1001  
 pclose subroutine 1002  
 performance monitor subroutines  
     pm\_delete\_program\_pgroup 1032  
     pm\_delete\_program\_pthread 1033  
     pm\_get\_data\_pgroup 1050  
     pm\_get\_data\_pgroup\_mx 1052  
     pm\_get\_data\_pthread 1054  
     pm\_get\_data\_pthread\_mx 1055  
     pm\_get\_program\_pgroup 1074  
     pm\_get\_program\_pgroup\_mx 1075  
     pm\_get\_program\_pthread 1077  
     pm\_get\_program\_pthread\_mx 1078  
     pm\_get\_tdata\_pgroup 1050  
     pm\_get\_Tdata\_pgroup 1050  
     pm\_get\_tdata\_pgroup\_mx 1052  
     pm\_get\_tdata\_pthread 1054  
     pm\_get\_Tdata\_pthread 1054  
     pm\_get\_tdata\_pthread\_mx 1055  
     pm\_initialize 1085  
     pm\_reset\_data\_pgroup 1091  
     pm\_reset\_data\_pthread 1093  
     pm\_set\_program\_pgroup 1109  
     pm\_set\_program\_pgroup\_mx 1111  
     pm\_set\_program\_pthread 1113  
     pm\_set\_program\_pthread\_mx 1115  
     pm\_start\_pgroup 1126  
     pm\_start\_pthread 1127  
     pm\_stop\_pgroup 1134  
     pm\_tstart\_pgroup subroutine 1126  
     pm\_tstart\_pthread subroutine 1127  
     pm\_tstop\_pgroup subroutine 1134  
 perfstat  
     perfstat\_partition\_total subroutine 1021  
     perfstat\_cpu subroutine 1003  
     perfstat\_cpu\_total subroutine 1004  
     perfstat\_disk subroutine 1006  
     perfstat\_disk\_total subroutine 1011  
     perfstat\_diskadapter subroutine 1008  
     perfstat\_diskpath subroutine 1009  
     perfstat\_memory\_total subroutine 1013  
     perfstat\_netbuffer subroutine 1014  
     perfstat\_netinterface subroutine 1015  
     perfstat\_netinterface\_total subroutine 1017  
     perfstat\_pagingspace subroutine 1018  
     perfstat\_partial\_reset subroutine 1019  
     perfstat\_partition\_total subroutine 1021  
     perfstat\_protocol subroutine 1022  
     perfstat\_reset subroutine 1023  
 permanent storage  
     writing file changes to 321  
 perror subroutine 1024  
 pglob parameter  
     freeing memory 485  
 physical volumes  
     querying 770  
 pipe subroutine 1025  
 pipes  
     closing 1002  
     creating 1025, 1144  
 plock subroutine 1026  
 pm\_delete\_program\_pgroup subroutine 1032  
 pm\_delete\_program\_pthread subroutine 1033

- pm\_get\_data\_pgroup subroutine 1050
- pm\_get\_data\_pgroup\_mx subroutine 1052
- pm\_get\_data\_pthread subroutine 1054
- pm\_get\_data\_pthread\_mx subroutine 1055
- pm\_get\_proctype 1060
- pm\_get\_program\_pgroup subroutine 1074
- pm\_get\_program\_pgroup\_mx subroutine 1075
- pm\_get\_program\_pthread subroutine 1077
- pm\_get\_program\_pthread\_mx subroutine 1078
- pm\_get\_tdata\_pgroup subroutine 1050
- pm\_get\_Tdata\_pgroup subroutine 1050
- pm\_get\_tdata\_pgroup\_mx subroutine 1052
- pm\_get\_tdata\_pthread subroutine 1054
- pm\_get\_Tdata\_pthread subroutine 1054
- pm\_get\_tdata\_pthread\_mx subroutine 1055
- pm\_initialize subroutine 1085
- pm\_reset\_data\_pgroup subroutine 1091
- pm\_reset\_data\_pthread subroutine 1093
- pm\_set\_program\_pgroup subroutine 1109
- pm\_set\_program\_pgroup\_mm subroutine 1111
- pm\_set\_program\_pgroup\_mx subroutine 1111
- pm\_set\_program\_pthread subroutine 1113
- pm\_set\_program\_pthread\_mm subroutine 1115
- pm\_set\_program\_pthread\_mx subroutine 1115
- pm\_start\_pgroup subroutine 1126
- pm\_start\_pthread subroutine 1127
- pm\_stop\_pgroup subroutine 1134
- pm\_tstart\_pgroup subroutine 1126
- pm\_tstart\_pthread subroutine 1127
- pm\_tstop\_pgroup subroutine 1134
- poll subroutine 1138
- pollset subroutines
  - pollset\_create 1141
  - pollset\_ctl 1141
  - pollset\_destroy 1141
  - pollset\_poll 1141
  - pollset\_query 1141
- pollset\_create subroutine 1141
- pollset\_ctl subroutine 1141
- pollset\_destroy subroutine 1141
- pollset\_poll subroutine 1141
- pollset\_query subroutine 1141
- popen subroutine 1144
- POSIX Realtime subroutines
  - posix\_fadvise 1145
  - posix\_fallocate 1146
  - posix\_madvise 1147
- POSIX SPAWN subroutines
  - posix\_spawn 1149
  - posix\_spawn\_file\_actions\_addclose 1153
  - posix\_spawn\_file\_actions\_adddup2 1154
  - posix\_spawn\_file\_actions\_addopen 1153
  - posix\_spawn\_file\_actions\_destroy 1155
  - posix\_spawn\_file\_actions\_init 1155
  - posix\_spawnattr\_destroy 1156
  - posix\_spawnattr\_getflags 1157
  - posix\_spawnattr\_getpgroup 1158
  - posix\_spawnattr\_getschedparam 1159
  - posix\_spawnattr\_getschedpolicy 1160
  - posix\_spawnattr\_getsigdefault 1161
  - posix\_spawnattr\_getsigmask 1162

- POSIX SPAWN subroutines (*continued*)
  - posix\_spawnattr\_init 1156
  - posix\_spawnattr\_setflags 1157
  - posix\_spawnattr\_setpgroup 1158
  - posix\_spawnattr\_setschedparam 1159
  - posix\_spawnattr\_setschedpolicy 1160
  - posix\_spawnattr\_setsigdefault 1161
  - posix\_spawnattr\_setsigmask 1162
  - posix\_spawnnp 1149
  - posix\_openpt Subroutine 1148
  - posix\_spawn subroutine 1149
  - posix\_spawn\_file\_actions\_addclose subroutine 1153
  - posix\_spawn\_file\_actions\_adddup2 subroutine 1154
  - posix\_spawn\_file\_actions\_addopen subroutine 1153
  - posix\_spawn\_file\_actions\_destroy subroutine 1155
  - posix\_spawn\_file\_actions\_init subroutine 1155
  - posix\_spawnattr\_destroy subroutine 1156
  - posix\_spawnattr\_getflags subroutine 1157
  - posix\_spawnattr\_getpgroup subroutine 1158
  - posix\_spawnattr\_getschedparam subroutine 1159
  - posix\_spawnattr\_getschedpolicy subroutine 1160
  - posix\_spawnattr\_getsigdefault subroutine 1161
  - posix\_spawnattr\_getsigmask subroutine 1162
  - posix\_spawnattr\_init subroutine 1156
  - posix\_spawnattr\_setflags subroutine 1157
  - posix\_spawnattr\_setpgroup subroutine 1158
  - posix\_spawnattr\_setschedparam subroutine 1159
  - posix\_spawnattr\_setschedpolicy subroutine 1160
  - posix\_spawnattr\_setsigdefault subroutine 1161
  - posix\_spawnattr\_setsigmask subroutine 1162
  - posix\_spawnnp subroutine 1149
  - posix\_trace\_getnext\_event subroutine 1163
  - posix\_trace\_timedgetnext\_event subroutine 1163
  - posix\_trace\_trygetnext\_event subroutine 1163
  - pow subroutine 785, 1166
  - power functions
    - computing 245
    - powf 1166
  - powf subroutine 1166
  - powl subroutine 1166
  - pre-editing space 556
  - print formatter subroutines
    - initialize 558
    - lineout 717
  - print lines
    - formatting 717
  - printer initialization 558
  - printf subroutine 1168
  - process accounting
    - displaying resource use 428
    - enabling and disabling 7
    - tracing process execution 1307
  - process credentials
    - reading 403
  - process environments
    - initializing run-time 581
    - reading 405
  - process group IDs
    - returning 370, 407
    - supplementary IDs
      - getting 383

- process group IDs *(continued)*
  - supplementary IDs *(continued)*
    - initializing 558
- process identification
  - alphanumeric user name 210
  - path name of controlling terminal 199
- process IDs
  - returning 407
- process initiation
  - creating child process 291
  - executing file 236
- process locks 1026
- process messages
  - getting message queue identifiers 881
  - providing control operations 879
  - reading from message queue 882
  - receiving from message queue 887
  - sending to message queue 885
- process priorities
  - getting or setting 412
  - returning scheduled priorities 411
- process program counters
  - histogram 1176
- process resource allocation
  - changing data space segments 122
  - controlling system consumption 424
  - getting size of descriptor table 363
  - locking into memory 1026
  - starting address sampling 1176
  - stopping address sampling 1176
- process resource use 428
- process signals
  - alarm 386
  - printing system signal messages 1182
  - sending to processes 580
- process subroutines (security and auditing)
  - getegid 370
  - geteuid 452
  - getgid 370
  - getgidx 370
  - getgroups 383
  - getpcred 403
  - getpenv 405
  - getuid 452
  - getuidx 452
  - initgroups 558
  - kleanup 581
- process user IDs
  - returning 452
- processes
  - closing pipes 1002
  - creating 291
  - getting process table entries 415
  - initializing run-time environment 581
  - initiating pipes 1144
  - suspending 982
  - terminating 2, 243, 580
  - tracing 1307
- processes subroutines
  - \_exit 243
  - abort 2

- processes subroutines *(continued)*
  - acct 7
  - atexit 243
  - brk 122
  - ctermid 199
  - cuserid 210
  - exec 236
  - exit 243
  - fork 291
  - getdtablesize 363
  - getpgrp 407
  - getpid 407
  - getppid 407
  - getpri 411
  - getpriority 412
  - getrlimit 424
  - getrlimit64 424
  - getrusage 428
  - getrusage64 428
  - kill 580
  - killpg 580
  - msgctl 879
  - msgget 881
  - msgrcv 882
  - msgsnd 885
  - msgxrcv 887
  - nice 412
  - pause 982
  - plock 1026
  - profil 1176
  - psignal 1182
  - ptrace 1307
  - sbrk 122
  - setpriority 412
  - setrlimit 424
  - setrlimit64 424
  - times 428
  - unatexit 243
  - vfork 291
  - vlimit 424
  - vtimes 428
- profil subroutine 1176
- program assertion
  - verifying 92
- proj\_execve subroutine 1178
- projdballoc subroutine 1179
- projdbfinit subroutine 1179
- projdbfree subroutine 1181
- psdanger subroutine 1181
- psignal subroutine 1182
- pthdb\_attr\_
  - pthdb\_attr\_addr 1185
  - pthdb\_attr\_detachstate 1185
  - pthdb\_attr\_guardsize 1185
  - pthdb\_attr\_inheritsched 1185
  - pthdb\_attr\_schedparam 1185
  - pthdb\_attr\_schedpolicy 1185
  - pthdb\_attr\_schedpriority 1185
  - pthdb\_attr\_scope 1185
  - pthdb\_attr\_stackaddr 1185
  - pthdb\_attr\_stacksize 1185

pthread\_attr\_ (continued)  
   pthread\_attr\_suspendstate 1185  
 pthread subroutines  
   pthread\_attr\_getinheritsched subroutine 1212  
   pthread\_attr\_getschedpolicy subroutine 1214  
   pthread\_attr\_setinheritsched subroutine 1212  
   pthread\_attr\_setschedpolicy subroutine 1214  
   pthread\_create\_withcred\_np 1245  
   pthread\_mutex\_timedlock 1272  
   pthread\_rwlock\_timedrdlock 1287  
   pthread\_rwlock\_timedwrlock 1289  
 pthread\_atfork subroutine 1208  
 pthread\_attr\_destroy subroutine 1209  
 pthread\_attr\_getdetachstate subroutine 1218  
 pthread\_attr\_getguardsize subroutine 1210  
 pthread\_attr\_getinheritsched subroutine 1212  
 pthread\_attr\_getschedparam subroutine 1213  
 pthread\_attr\_getschedpolicy subroutine 1214  
 pthread\_attr\_getscope subroutine 1219  
 pthread\_attr\_getstackaddr subroutine 1215  
 pthread\_attr\_getstacksize subroutine 1216  
 pthread\_attr\_getukeyset\_np 1221  
 pthread\_attr\_init subroutine 1217  
 pthread\_attr\_setdetachstate subroutine 1218  
 pthread\_attr\_setguardsize subroutine 1210  
 pthread\_attr\_setinheritsched subroutine 1212  
 pthread\_attr\_setschedparam subroutine 1222  
 pthread\_attr\_setschedpolicy subroutine 1214  
 pthread\_attr\_setscope subroutine 1219  
 pthread\_attr\_setstackaddr subroutine 1223  
 pthread\_attr\_setstacksize subroutine 1224  
 pthread\_attr\_setsuspendstate\_np and  
   pthread\_attr\_getsuspendstate\_np subroutine 1225  
 pthread\_attr\_setukeyset\_np 1221  
 pthread\_cancel subroutine 1231  
 pthread\_cleanup\_pop subroutine 1232  
 pthread\_cleanup\_push subroutine 1232  
 pthread\_cond\_broadcast subroutine 1235  
 pthread\_cond\_destroy subroutine 1233  
 PTHREAD\_COND\_INITIALIZER macro 1235  
 pthread\_cond\_signal subroutine 1235  
 pthread\_cond\_timedwait subroutine 1236  
 pthread\_cond\_wait subroutine 1236  
 pthread\_condattr\_destroy subroutine 1238  
 pthread\_condattr\_getclock subroutine 1240  
 pthread\_condattr\_getpshared subroutine 1241  
 pthread\_condattr\_setclock subroutine 1240  
 pthread\_condattr\_setpshared subroutine 1242  
 pthread\_create subroutine 1243  
 pthread\_create\_withcred\_np subroutine 1245  
 pthread\_delay\_np subroutine 1246  
 pthread\_equal subroutine 1247  
 pthread\_exit subroutine 1248  
 pthread\_get\_expiration\_np subroutine 1250  
 pthread\_getconcurrency subroutine 1251  
 pthread\_getcpuclockid subroutine 1252  
 pthread\_getrusage\_np subroutine 1252  
 pthread\_getschedparam subroutine 1255  
 pthread\_getspecific subroutine 1256  
 pthread\_getunique\_np subroutine 1261  
 pthread\_join subroutine 1262  
 pthread\_key\_create subroutine 1263  
 pthread\_key\_delete subroutine 1264  
 pthread\_kill subroutine 1265  
 pthread\_lock\_global\_np subroutine 1266  
 pthread\_mutex\_destroy subroutine 1267  
 pthread\_mutex\_init subroutine 1267  
 PTHREAD\_MUTEX\_INITIALIZER macro 1270  
 pthread\_mutex\_lock subroutine 1270  
 pthread\_mutex\_timedlock subroutine 1272  
 pthread\_mutex\_trylock subroutine 1270  
 pthread\_mutexattr\_destroy subroutine 1273  
 pthread\_mutexattr\_getkind\_np subroutine 1275  
 pthread\_mutexattr\_gettype subroutine 1280  
 pthread\_mutexattr\_init subroutine 1273  
 pthread\_mutexattr\_setkind\_np subroutine 1281  
 pthread\_mutexattr\_settype subroutine 1280  
 pthread\_once subroutine 1283  
 PTHREAD\_ONCE\_INIT macro 1284  
 pthread\_rwlock\_timedrdlock subroutine 1287  
 pthread\_rwlock\_timedwrlock subroutine 1289  
 pthread\_self subroutine 1295  
 pthread\_setcancelstate subroutine 1296  
 pthread\_setschedparam subroutine 1297  
 pthread\_setschedprio subroutine 1299  
 pthread\_setspecific subroutine 1256  
 pthread\_signal\_to\_cancel\_np subroutine 1300  
 pthread\_spin\_destroy subroutine 1301  
 pthread\_spin\_init subroutine 1301  
 pthread\_suspend\_np, pthread\_unsuspend\_np and  
   pthread\_continue\_np subroutine 1304  
 pthread\_unlock\_global\_np subroutine 1305  
 pthread\_yield subroutine 1306  
 pthreads subroutines  
   posix\_trace\_getnext\_event subroutine 1163  
   posix\_trace\_timedgetnext\_event subroutine 1163  
   posix\_trace\_trygetnext\_event subroutine 1163  
   pthread\_setschedprio subroutine 1299  
 ptrace subroutine 1307  
 ptracex subroutine 1307  
 ptsname subroutine 1320  
 putc subroutine 1321  
 putc\_unlocked subroutine 349  
 putchar subroutine 1321  
 putchar\_unlocked subroutine 349  
 putconfattr subroutine 350  
 putconfattr subroutine 1323  
 putenv subroutine 1326  
 putgrent subroutine 1326  
 putgroupattr subroutine 375  
 putgroupattr subroutine 1327  
 putgprattr Subroutine 384  
 putportattr Subroutine 408  
 putpwent subroutine 422  
 putroleattr Subroutine 431  
 puts subroutine 1331  
 puttcattr subroutine 440  
 putuserattr subroutine 454  
 putuserattr subroutine 1333  
 putuserpw subroutine 468  
 putuserpwhist subroutine 468  
 putuserpwx subroutine 1337

putusraclattr Subroutine 472  
pututline subroutine 474  
putw subroutine 1321  
putwc subroutine 1339  
putwchar subroutine 1339  
putws subroutine 1341

## Q

queues  
  inserting and removing elements 559  
quotient and remainder  
  imaxdiv 535

## R

radix-independent exponents  
  logbf 746, 748  
  logbl 746, 748  
read operations  
  asynchronous 50  
  binary files 310  
read-write file pointers  
  moving 765  
readdir subroutine 943  
readdir64 subroutine 943  
real floating types  
  fpclassify 309  
real value subroutines  
  creal 190  
  crealf 190  
  creall 190  
realloc subroutine 778  
regular expressions  
  matching patterns 177  
remque subroutine 559  
resabs subroutine 386  
reset\_speed subroutine 342  
resinc subroutine 386  
resource information 1252  
resources subroutines  
  pthread\_getrusage\_np 1252  
restimer subroutine 446  
rewind subroutine 318  
rewinddir subroutine 943  
rewinddir64 subroutine 943  
rounding direction  
  fegetround 267  
  fesetround 267  
rounding numbers  
  llrint 725  
  llrintf 725  
  llrintl 725  
  llround 726  
  llroundf 726  
  llroundl 726  
  lrint 761  
  lrintd128 761  
  lrintd32 761  
  lrintd64 761  
  lrintf 761

rounding numbers (*continued*)  
  lrintl 761  
  lround 763  
  lroundf 763  
  lroundl 763  
rpc file  
  handling 427  
rpow subroutine 785  
run-time environment  
  initializing 581

## S

sbrk subroutine 122  
sdiv subroutine 785  
security library subroutines  
  authentocatex 115  
  chpassx 156  
  getconfattrs 355  
  getgroupattrs 378  
  getuserattrs 460  
  getuserpwx 470  
  loginrestrictionsx 754  
  newpassx 903  
  passwdexpiredx 974  
  putconfattrs 1323  
  putgroupattrs 1327  
  putuserattrs 1333  
  putuserpwx 1337  
security subroutines  
  getuinfox 453  
seed48 subroutine 221  
seekdir subroutine 943  
seekdir64 subroutine 943  
set\_speed subroutine 342  
setfsent subroutine 369  
setfsent\_r subroutine 435  
setgrent subroutine 371  
setitimer subroutine 386  
setkey subroutine 191  
setpriority subroutine 412  
setpwent subroutine 422  
setrlimit subroutine 424  
setrlimit64 subroutine 424  
setrpcent subroutine 427  
setsockopt subroutine 529  
settimeofday subroutine 445  
settimer subroutine 446  
setttyent subroutine 450  
setutent subroutine 474  
setvfsent subroutine 476  
shell command-line flags 397  
SIGALRM signal 388  
SIGIOT signal 2  
signal names  
  formatting 1182  
sine functions  
  csin 193  
  csinf 193  
  csinl 193  
single-byte to wide-character conversion 124



- SJIS character conversions 575
- sjtojis subroutine 575
- sjtojuj subroutine 575
- snprintf subroutine 1168
- socket options
  - setting 529
- sockets kernel service subroutines
  - setsockopt 529
- sockets network library subroutines
  - endhostent 911
  - gethostent 910
  - inet\_aton 557
- special files
  - creating 837
- sprintf subroutine 1168
- square root subroutines
  - csqrt 195
  - csqrtf 195
  - csqrtl 195
- srand48 subroutine 221
- SRC subroutines
  - addssys 33
  - chssys 162
  - delssys 212
  - getssys 437
- SRC subsys record
  - adding 33
- SRC subsys structure
  - initializing 211
- Statistics subroutines
  - perfstat\_cpu 1003
  - perfstat\_cpu\_total 1004
  - perfstat\_disk 1006
  - perfstat\_disk\_total 1011
  - perfstat\_diskadapter 1008
  - perfstat\_diskpath 1009
  - perfstat\_memory\_total 1013
  - perfstat\_netbuffer 1014
  - perfstat\_netinterface 1015
  - perfstat\_netinterface\_total 1017
  - perfstat\_pagingspace 1018
  - perfstat\_protocol 1022
  - perfstat\_reset 1023
- status indicators
  - beeping 539
  - drawing 543
  - hiding 544
- step subroutine 177
- stime subroutine 446
- streams
  - checking status 270
  - closing 253
  - flushing 253
  - opening 288
  - repositioning file pointers 318
  - writing to 253
- string conversion
  - long integers to base-64 ASCII 1
- string manipulation subroutines
  - advance 177
  - bcmp 118
- string manipulation subroutines (*continued*)
  - bcopy 118
  - bzero 118
  - compile 177
  - ffs 118
  - fgets 434
  - fnmatch 286
  - fputs 1331
  - gets 434
  - puts 1331
  - step 177
- strings
  - bit string operations 118
  - byte string operations 118
  - copying 118
  - drawing text strings 555
  - matching against pattern parameters 286
  - reading bytes into arrays 434
  - writing to standard output streams 1331
  - zeroing out 118
- subroutine
  - pcap\_close 982
  - pcap\_compile 983
  - pcap\_datalink 983
  - pcap\_dispatch 984
  - pcap\_dump 985
  - pcap\_dump\_close 986
  - pcap\_dump\_open 987
  - pcap\_file 988
  - pcap\_fileno 988
  - pcap\_geterr 989
  - pcap\_is\_swapped 990
  - pcap\_lookupdev 990
  - pcap\_lookupnet 991
  - pcap\_loop 992
  - pcap\_major\_version 993
  - pcap\_next 994
  - pcap\_open\_live 995
  - pcap\_open\_offline 996
  - pcap\_perror 998
  - pcap\_setfilter 999
  - pcap\_snapshot 1000
  - pcap\_stats 1000
  - pcap\_strerror 1001
- subroutines
  - LAPI\_Addr\_get 588
  - LAPI\_Addr\_set 589
  - LAPI\_Address 591
  - LAPI\_Address\_init 592
  - LAPI\_Address\_init64 594
  - LAPI\_Amsend 596
  - LAPI\_Amsendv 601
  - LAPI\_Fence 609
  - LAPI\_Get 610
  - LAPI\_Getcptr 613
  - LAPI\_Getv 614
  - LAPI\_Gfence 618
  - LAPI\_Init 619
  - LAPI\_Msg\_string 624
  - LAPI\_Msgpoll 626
  - LAPI\_Nopoll\_wait 628

- subroutines (*continued*)
  - LAPI\_Probe 629
  - LAPI\_Purge\_totask 630
  - LAPI\_Put 632
  - LAPI\_Putv 634
  - LAPI\_Qenv 638
  - LAPI\_Resume\_totask 641
  - LAPI\_Rmw 642
  - LAPI\_Rmw64 646
  - LAPI\_Senv 650
  - LAPI\_Setcptr 652
  - LAPI\_Setcptr\_wstatus 654
  - LAPI\_Term 655
  - LAPI\_Util 657
  - LAPI\_Waitcptr 668
  - LAPI\_Xfer 670
- Subroutines
  - perfstat\_cpu 1003
  - perfstat\_cpu\_total 1004
  - perfstat\_disk\_total 1006, 1011
  - perfstat\_diskpath 1009
  - perfstat\_memory\_total 1013
  - perfstat\_netinterface\_total 1015, 1017
- subsystem objects
  - modifying 162
  - removing 212
- subsystem records
  - reading 437, 439
- supplementary process group IDs
  - getting 383
  - initializing 558
- swapcontext Subroutine 788
- swprintf subroutine 327
- swscanf subroutine 331
- symbol-handling subroutine
  - knlist 582
- symbols
  - translating names to addresses 582
- sys\_siglist vector 1182
- SYSP\_V\_IOSTRUN
  - sys\_parm 1006, 1008, 1009
- system auditing 98
- system data objects
  - auditing modes 105
- system event audits
  - getting or setting status 102
- system resources
  - setting maximums 424
- system signal messages 1182
- system variables
  - determining values 181

## T

- tellmdir subroutine 943
- tellmdir64 subroutine 943
- terminal baud rate
  - get 342
  - set 342
- text area
  - hiding 555

- text locks 1026
- text strings
  - drawing 555
- Thread-safe C Library
  - subroutines
    - 164\_r 587
- Thread-Safe C Library 373, 374, 435
  - subroutines
    - getfsent\_r 435
    - getlogin\_r 395
    - getsfile\_r 435
    - setfsent\_r 435
- threads
  - getting thread table entries 443
- Threads Library 1297
  - condition variables
    - creation and destruction 1233, 1235
    - creation attributes 1238, 1241, 1242
    - signalling a condition 1235
    - waiting for a condition 1236
  - DCE compatibility subroutines
    - pthread\_delay\_np 1246
    - pthread\_get\_expiration\_np 1250
    - pthread\_getunique\_np 1261
    - pthread\_lock\_global\_np 1266
    - pthread\_mutexattr\_getkind\_np 1275
    - pthread\_mutexattr\_setkind\_np 1281
    - pthread\_signal\_to\_cancel\_np 1300
    - pthread\_unlock\_global\_np 1305
  - mutexes
    - creation and destruction 1270
    - creation attributes 1280
    - locking 1270
    - pthread\_mutexattr\_destroy 1273
    - pthread\_mutexattr\_init 1273
  - process creation
    - pthread\_atfork subroutine 1208
  - pthread\_attr\_getguardsize subroutine 1210
  - pthread\_attr\_setguardsize subroutine 1210
  - pthread\_getconcurrency subroutine 1251
  - pthread\_mutex\_destroy 1267
  - pthread\_mutex\_init 1267
- scheduling
  - dynamic thread control 1255, 1306
  - thread creation attributes 1213, 1222
- signal, sleep, and timer handling
  - pthread\_kill subroutine 1265
- thread-specific data
  - pthread\_getspecific subroutine 1256
  - pthread\_key\_create subroutine 1263
  - pthread\_key\_delete subroutine 1264
  - pthread\_setspecific subroutine 1256
- threads
  - cancellation 1231, 1296
  - creation 1243
  - creation attributes 1209, 1215, 1216, 1217, 1218, 1219, 1223, 1224, 1225, 1304
  - ID handling 1247, 1295
  - initialization 1283, 1284
  - termination 1232, 1248, 1262

- time
  - displaying and setting 445
  - reporting used CPU time 169
  - synchronizing system clocks 35
- time format conversions 200
- time manipulation subroutines
  - absinterval 386
  - adjtime 35
  - alarm 386
  - asctime 200
  - clock 169
  - clock\_getres 170
  - clock\_gettime 170
  - clock\_settime 170
  - ctime 200
  - difftime 200
  - ftime 445
  - getinterval 386
  - getitimer 386
  - gettimeofday 445
  - gettimer 446
  - gettimerid 449
  - gmtime 200
  - incinterval 386
  - localtime 200
  - mktime 200
  - resabs 386
  - resinc 386
  - restimer 446
  - setitimer 386
  - settimeofday 445
  - settimer 446
  - stime 446
  - time 446
  - tzset 200
  - ualarm 386
- time subroutine 446
- time subroutines
  - asctime64 202
  - asctime64\_r 204
  - ctime64 202
  - ctime64\_r 204
  - difftime64 202
  - gmtime64 202
  - gmtime64\_r 204
  - localtime64 202
  - localtime64\_r 204
  - mktime64 202
- timer
  - getting or setting values 446
- timer subroutines
  - clock\_getcpuclockid 169
  - clock\_nanosleep 172
  - pthread\_condattr\_getclock 1240
  - pthread\_condattr\_setclock 1240
  - pthread\_getcpuclockid 1252
- times subroutine 428
- toascii subroutine 183
- tojhira subroutine 576
- tojkata subroutine 576
- tolower subroutine 576

- toupper subroutine 576
- tolower subroutine 183
- toujis subroutine 576
- toupper subroutine 183
- trace
  - install\_lwcf\_handler subroutine 560
  - mt\_trce subroutine 891
- transforming text 692
- trunc subroutine 277
- trusted processes
  - initializing run-time environment 581
- tty description file
  - querying 450
- tty subroutines
  - endttyent 450
  - getttyent 450
  - getttynam 450
  - setttyent 450
- tzset subroutine 200

## U

- ualarm subroutine 386
- uitrunc subroutine 277
- UJIS character conversions 575
- ujtojis subroutine 575
- ujtosj subroutine 575
- umul\_dbl subroutine 3
- unatexit subroutine 243
- unbiased exponents
  - ilogbf 534
  - ilogbl 534
- unordered subroutine 167
- user accounts
  - checking validity 164
- user authentication data
  - accessing 468
- user database
  - accessing group information 371, 375
  - accessing user information 350, 422, 454
- user information
  - accessing 350, 422, 454
  - accessing group information 371, 375
  - searching buffer 453
- user login name
  - getting 393
- users
  - authenticating 166
- utmpname subroutine 474

## V

- vectors
  - sys\_siglist 1182
- vfork subroutine 291
- vfprintf subroutine 1168
- VFS (Virtual File System)
  - getting file entries 476
  - returning mount status 847
- virtual memory
  - mapping file-system objects 843

- vlimit subroutine 424
- volume groups
  - querying 774
  - querying all varied on-line 777
- vprintf subroutine 1168
- vsprintf subroutine 1168
- vtimes subroutine 428
- vwsprintf subroutine 1168

## W

- wide character subroutines
  - fgetwc 477
  - fgetws 480
  - fputwc 1339
  - fputws 1341
  - getwc 477
  - getwchar 477
  - getws 480
  - is\_wctype 574
  - iswalnum 572
  - iswalpha 572
  - iswcntrl 572
  - iswctype subroutine 574
  - iswdigit 572
  - iswgraph 572
  - iswlower 572
  - iswprint 572
  - iswpunct 572
  - iswspace 572
  - iswupper 572
  - iswxdigit 572
  - putwc 1339
  - putwchar 1339
  - putws 1341
- wide characters
  - checking character class 572
  - converting
    - from multibyte 804, 805
  - determining properties 574
  - reading from input stream 477, 480
  - writing to output stream 1339, 1341
- words
  - returning from input streams 347
- wprintf subroutine 327
- write operations
  - asynchronous 61
  - binary files 310
- wscanf subroutine 331
- wsprintf subroutine 1168

## Y

- y0 subroutine 119
- y1 subroutine 119
- yn subroutine 119

---

# Readers' Comments — We'd Like to Hear from You

**AIX 5L Version 5.3**

**Technical Reference: Base Operating System and Extensions, Volume 1**

**Publication No. SC23-4913-04**

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- Send your comments via e-mail to: [aix6koub@austin.ibm.com](mailto:aix6koub@austin.ibm.com)

If you would like a response from IBM, please fill in the following information:

\_\_\_\_\_

Name

\_\_\_\_\_

Address

\_\_\_\_\_

Company or Organization

\_\_\_\_\_

Phone No.

\_\_\_\_\_

E-mail address



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation  
Information Development  
Department 04XA-905-6C006  
11501 Burnet Road  
Austin, TX 78758-3493



Fold and Tape

Please do not staple

Fold and Tape





Printed in U.S.A.

SC23-4913-04

