



man pages section 7: Device and Network Interfaces

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 816-5177-10
January 2005

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



050105@10536



Contents

Preface 11

Introduction 17

Intro(7) 18

Device and Network Interfaces 21

aac(7D) 22

adp(7D) 23

adpu320(7D) 25

afb(7d) 27

amr(7D) 28

arp(7P) 29

asy(7D) 33

ata(7D) 36

audio(7I) 40

audio810(7D) 51

audiocs(7D) 54

audioens(7D) 57

audio_support(7I) 61

audiots(7D) 63

av1394(7D) 66

bbc_beep(7D) 67

bd(7M) 68

bge(7D) 70

bpp(7D) 74

bscv(7D) 79

bufmod(7M) 80
cadp160(7D) 84
cadp(7D) 85
cdio(7I) 89
ce(7D) 97
cgsix(7D) 101
chs(7D) 102
cmdk(7D) 103
connld(7M) 104
console(7D) 106
cpr(7) 107
cpuid(7D) 109
ctfs(7FS) 111
ctsmc(7D) 112
cvc(7D) 113
cvcredir(7D) 114
dad(7D) 115
daplt(7D) 118
dbri(7D) 119
devfs(7FS) 124
devinfo(7D) 125
dkio(7I) 126
dlcosmk(7ipp) 136
dlpi(7P) 137
dmfe(7D) 140
dnet(7D) 142
dpt(7D) 144
dr(7d) 146
dscpmk(7ipp) 147
dtrace(7D) 148
e1000g(7D) 149
ecpp(7D) 154
ehci(7D) 160
elxl(7D) 162
eri(7D) 165
esp(7D) 169
fas(7D) 175
fasttrap(7D) 183

fbio(7I) 184
fbt(7D) 186
fcip(7D) 187
fcp(7D) 190
fctl(7D) 191
fd(7D) 192
fdio(7I) 197
ffb(7D) 201
flowacct(7ipp) 202
fp(7d) 203
FSS(7) 205
ge(7D) 208
gld(7D) 212
glm(7D) 222
gpio_87317(7D) 227
grbeep(7d) 228
hci1394(7D) 229
hdio(7I) 230
hid(7D) 232
hme(7D) 234
hpfc(7D) 239
hsfs(7FS) 241
hubd(7D) 244
i2o_bs(7D) 246
i2o_scsi(7D) 248
ib(7D) 249
ibcm(7D) 252
ibd(7D) 253
ibdm(7D) 256
ibmf(7) 257
ibtl(7D) 258
icmp6(7P) 259
icmp(7P) 261
idn(7d) 263
ifb(7d) 266
ifp(7D) 267
if_tcp(7P) 271
inet6(7P) 280

inet(7P) 285
ip6(7P) 288
ip(7P) 295
ipgpc(7ipp) 301
ipqos(7ipp) 303
iprb(7D) 305
ipsec(7P) 308
ipsecah(7P) 312
ipsecesp(7P) 313
isdnio(7I) 314
isp(7D) 327
ixgb(7d) 333
jfb(7D) 335
jfca(7D) 336
kb(7M) 337
kdmouse(7D) 346
kmdb(7d) 347
kstat(7D) 348
ksyms(7D) 349
ldterm(7M) 351
llc1(7D) 354
llc2(7D) 357
lockstat(7D) 363
lofi(7D) 364
lofs(7FS) 365
log(7D) 366
logi(7D) 370
lp(7D) 371
m64(7D) 373
md(7D) 374
mediator(7D) 378
mem(7D) 381
mhd(7i) 383
mixer(7I) 388
mpt(7D) 398
msglog(7D) 402
msm(7D) 403
mt(7D) 404

mtio(7I) 405
ncrs(7D) 417
null(7D) 424
objfs(7FS) 425
ocf_ibutton(7D) 426
ohci(7D) 427
openprom(7D) 429
pcata(7D) 434
pcelx(7D) 436
pcf_s(7FS) 438
pcic(7D) 443
pckt(7M) 445
pcmem(7D) 446
pcn(7D) 447
pcram(7D) 449
pcscsi(7D) 452
pcser(7D) 453
pfb(7D) 455
pf_key(7P) 456
pfmod(7M) 466
pipemod(7M) 469
pm(7D) 470
poll(7d) 474
prnio(7I) 479
profile(7D) 483
ptem(7M) 484
ptm(7D) 485
pts(7D) 487
pty(7D) 489
qfe(7d) 492
qlc(7D) 496
quotactl(7I) 498
qus(7D) 500
ramdisk(7D) 507
random(7D) 509
rarp(7P) 511
route(7P) 512
routing(7P) 516

rtls(7D) 518
sad(7D) 520
sbpro(7D) 523
scmi2c(7d) 528
scsa1394(7D) 529
scsa2usb(7D) 532
scsi_vhci(7D) 537
sctp(7P) 539
sd(7D) 543
sdt(7D) 549
se(7D) 550
se_hdlc(7D) 554
ses(7D) 557
sesio(7I) 559
sf(7D) 560
sgen(7D) 562
sip(7P) 567
sk98sol(7D) 568
skfp(7D) 574
slp(7P) 576
smbus(7D) 578
social(7D) 579
sockio(7I) 581
sppptun(7M) 582
spwr(7D) 583
ssd(7D) 584
st(7D) 588
stp4020(7D) 603
streamio(7I) 604
su(7D) 620
sxp(7D) 623
symhisl(7D) 626
sysmsg(7D) 629
systrace(7D) 630
tavor(7D) 631
tcp(7P) 633
termio(7I) 638
termiox(7I) 659

ticlts(7D) 665
timod(7M) 667
tirdwr(7M) 669
tmpfs(7FS) 671
tokenmt(7ipp) 673
tpf(7D) 676
tsalarm(7D) 677
tswtclmt(7ipp) 680
ttcompat(7M) 681
tty(7D) 689
ttymux(7D) 690
tun(7M) 691
uata(7D) 695
udfs(7FS) 698
udp(7P) 699
ufs(7FS) 701
ugen(7D) 704
uhci(7D) 720
usba(7D) 721
usb_ac(7D) 723
usb_ah(7M) 726
usb_as(7D) 727
usbbkm(7M) 729
usb_mid(7D) 731
usbms(7M) 732
usbprn(7D) 735
usbser_edge(7D) 740
uscsi(7I) 743
usoc(7D) 747
visual_io(7I) 749
vni(7d) 755
volfs(7FS) 756
vuidmice(7M) 758
wrsm(7D) 762
wrsm(7D) 764
wscons(7D) 765
xge(7D) 774
xmemfs(7FS) 776

zcons(7D) 778
zero(7D) 779
zs(7D) 780
zsh(7D) 783
zulu(7d) 787

Index 789

Preface

Both novice users and those familiar with the SunOS operating system can use online man pages to obtain information about the system and its features. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

Overview

The following contains a brief description of each man page section and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character-set tables.
- Section 6 contains available games and demos.
- Section 7 describes various special files that refer to specific hardware peripherals and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.

- Section 9 provides reference information needed to write device drivers in the kernel environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer can include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

NAME	This section gives the names of the commands or functions documented, followed by a brief description of what they do.
SYNOPSIS	This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full path name is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.
	The following special characters are used in this section:
[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.
. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename ...".
	Separator. Only one of the arguments separated by this character can be specified at a time.
{ }	Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.

PROTOCOL	This section occurs only in subsection 3R to indicate the protocol description file.
DESCRIPTION	This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES. Interactive commands, subcommands, requests, macros, and functions are described under USAGE.
IOCTL	This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> .
OPTIONS	This section lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.
OPERANDS	This section lists the command operands and describes how they affect the actions of the command.
OUTPUT	This section describes the output – standard output, standard error, or output files – generated by the command.
RETURN VALUES	If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.
ERRORS	On failure, most functions place an error code in the global variable <code>errno</code> indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the

	<p>conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.</p>
USAGE	<p>This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality:</p> <ul style="list-style-type: none"> Commands Modifiers Variables Expressions Input Grammar
EXAMPLES	<p>This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command-line entry and machine response is shown. Whenever an example is given, the prompt is shown as <code>example%</code>, or if the user must be superuser, <code>example#</code>. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.</p>
ENVIRONMENT VARIABLES	<p>This section lists any environment variables that the command or function affects, followed by a brief description of the effect.</p>
EXIT STATUS	<p>This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion, and values other than zero for various error conditions.</p>
FILES	<p>This section lists all file names referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.</p>
ATTRIBUTES	<p>This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <code>attributes(5)</code> for more information.</p>
SEE ALSO	<p>This section lists references to other man pages, in-house documentation, and outside publications.</p>

DIAGNOSTICS	This section lists diagnostic messages with a brief explanation of the condition causing the error.
WARNINGS	This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.
NOTES	This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.
BUGS	This section describes known bugs and, wherever possible, suggests workarounds.

Introduction

Intro(7)

NAME	Intro – introduction to special files
DESCRIPTION	<p>This section describes various device and network interfaces available on the system. The types of interfaces described include character and block devices, STREAMS modules, network protocols, file systems, and ioctl requests for driver subsystems and classes.</p> <p>This section contains the following major collections:</p> <p>(7D) The system provides drivers for a variety of hardware devices, such as disk, magnetic tapes, serial communication lines, mice, and frame buffers, as well as virtual devices such as pseudo-terminals and windows.</p> <p>This section describes special files that refer to specific hardware peripherals and device drivers. STREAMS device drivers are also described. Characteristics of both the hardware device and the corresponding device driver are discussed where applicable.</p> <p>An application accesses a device through that device's special file. This section specifies the device special file to be used to access the device as well as application programming interface (API) information relevant to the use of the device driver.</p> <p>All device special files are located under the <code>/devices</code> directory. The <code>/devices</code> directory hierarchy attempts to mirror the hierarchy of system busses, controllers, and devices configured on the system. Logical device names for special files in <code>/devices</code> are located under the <code>/dev</code> directory. Although not every special file under <code>/devices</code> will have a corresponding logical entry under <code>/dev</code>, whenever possible, an application should reference a device using the logical name for the device. Logical device names are listed in the FILES section of the page for the device in question.</p> <p>This section also describes driver configuration where applicable. Many device drivers have a driver configuration file of the form <code>driver_name.conf</code> associated with them (see <code>driver.conf(4)</code>). The configuration information stored in the driver configuration file is used to configure the driver and the device. Driver configuration files are located in <code>/kernel/drv</code> and <code>/usr/kernel/drv</code>. Driver configuration files for platform dependent drivers are located in <code>/platform/`uname -i`/kernel/drv</code> where <code>`uname -i`</code> is the output of the <code>uname(1)</code> command with the <code>-i</code> option.</p> <p>Some driver configuration files may contain user configurable properties. Changes in a driver's configuration file will not take effect until the system is rebooted or the driver has been removed and re-added (see <code>rem_drv(1M)</code> and <code>add_drv(1M)</code>).</p> <p>(7FS) This section describes the programmatic interface for several file systems supported by SunOS.</p>

- (7I) This section describes ioctl requests which apply to a class of drivers or subsystems. For example, ioctl requests which apply to most tape devices are discussed in [mtio\(7I\)](#). Ioctl requests relevant to only a specific device are described on the man page for that device. The page for the device in question should still be examined for exceptions to the ioctls listed in section 7I.
- (7M) This section describes STREAMS modules. Note that STREAMS drivers are discussed in section 7D. [streamio\(7I\)](#) contains a list of ioctl requests used to manipulate STREAMS modules and interface with the STREAMS framework. Ioctl requests specific to a STREAMS module will be discussed on the man page for that module.
- (7P) This section describes various network protocols available in SunOS.
- SunOS supports both socket-based and STREAMS-based network communications. The Internet protocol family, described in [inet\(7P\)](#), is the primary protocol family supported by SunOS, although the system can support a number of others. The raw interface provides low-level services, such as packet fragmentation and reassembly, routing, addressing, and basic transport for socket-based implementations. Facilities for communicating using an Internet-family protocol are generally accessed by specifying the AF_INET address family when binding a socket; see [socket\(3SOCKET\)](#) for details.
- Major protocols in the Internet family include:
- The Internet Protocol (IP) itself, which supports the universal datagram format, as described in [ip\(7P\)](#). This is the default protocol for SOCK_RAW type sockets within the AF_INET domain.
 - The Transmission Control Protocol (TCP); see [tcp\(7P\)](#). This is the default protocol for SOCK_STREAM type sockets.
 - The User Datagram Protocol (UDP); see [udp\(7P\)](#). This is the default protocol for SOCK_DGRAM type sockets.
 - The Address Resolution Protocol (ARP); see [arp\(7P\)](#).
 - The Internet Control Message Protocol (ICMP); see [icmp\(7P\)](#).

SEE ALSO [add_drv\(1M\)](#), [rem_drv\(1M\)](#), [intro\(3\)](#), [ioctl\(2\)](#), [socket\(3SOCKET\)](#), [driver.conf\(4\)](#), [arp\(7P\)](#), [icmp\(7P\)](#), [inet\(7P\)](#), [ip\(7P\)](#), [mtio\(7I\)](#), [st\(7D\)](#), [streamio\(7I\)](#), [tcp\(7P\)](#), [udp\(7P\)](#)

System Administration Guide: IP Services

STREAMS Programming Guide

Writing Device Drivers

Intro(7)

Device and Network Interfaces

aac(7D)

NAME aac – SCSI HBA driver for Adaptec AdvancedRAID Controller

DESCRIPTION The aac plain SCSI host bus adapter driver is a SCSI-compliant nexus driver that supports the Adaptec 2200S/ 2120S SCSI RAID card and Dell PERC 3Di SCSI RAID controller.

The aac driver ports from FreeBSD and only supports basic RAID disk I/O functions.

DRIVER CONFIGURATION There are no user configurable parameters available. Please configure your hardware through BIOS.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	x86

FILES /kernel/drv/aac
32-bit ELF kernel module

/kernel/drv/aac.conf
Configuration file. (Contains no user-configurable options).

SEE ALSO `prtconf(1M)`, `attributes(5)`, `scsi_hba_attach(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Small Computer System Interface-2 (SCSI-2)

NAME	adp – Low-level module for controllers based on Adaptec AIC-7870P and AIC-7880P SCSI chips
DESCRIPTION	<p>The adp module provides low-level interface routines between the common disk/tape I/O system and SCSI (Small Computer System Interface) controllers based on the Adaptec AIC-7870P and AIC-7880P SCSI chips. These controllers include the Adaptec AHA-2940, AHA-2940W, AHA-2940U, AHA-2940UW, AHA-3940, and AHA-3940W, as well as motherboards with embedded AIC-7870P and AIC-7880P SCSI chips.</p> <p>Supported devices are AIC-7850, AIC-7860, AIC-7870, AIC-7880 and AIC-7895.</p> <p>The adp module can be configured for disk and streaming tape support for one or more host adapter boards, each of which must be the sole initiator on a SCSI bus. Auto-configuration code determines if the adapter is present at the configured address and what types of devices are attached to the adapter.</p>
PRECONFIGURE	The Plug N Play SCAM Support option is not supported.
Known Problems and Limitations	<ul style="list-style-type: none"> ■ To use the AHA-3940 or AHA-3940W adapters, the motherboard must have a BIOS that supports the DEC PCI-to-PCI Bridge chip on the host bus adapter. ■ User-level programs have exhibited problems on some PCI systems with an Adaptec AHA-2940 or AHA-2940W card and certain motherboards. If problems with user-level programs occur, use the BIOS setup to disable write-back CPU caching (or all caching if there is no control over the caching algorithm). The affected motherboards are: <ul style="list-style-type: none"> - PCI motherboards with a 60-MHz Pentium chip, with PCI chipset numbers S82433LX Z852 and S82434LX Z850. The part numbers of the Intel motherboards are AA616393-007 and AA615988-009. - PCI motherboards with a 90-MHz Pentium chip, with PCI chipset numbers S82433NX Z895, S82434NX Z895, and S82434NX Z896. The part number of the Intel motherboard is 541286-005. (Some Gateway 2000 systems use this motherboard.) - AA-619772-002 motherboard with 82433LX Z852 and 82434LX Z882 chips causes random memory inconsistencies. Return the motherboard to the vendor for a replacement. ■ If the AHA-2940 SCSI adapter does not recognize the Quantum Empire 1080S, HP 3323 SE or other SCSI disk drive, reduce the Synchronous Transfer rate on the Adaptec controller to 8 Mbps. ■ The AHA-3940 has been certified by Adaptec to work on specific systems; however, some testing has shown that the Solaris operating environment works properly in some of those systems and not in others.
CONFIGURATION	<p>Use the Adaptec configuration utility to perform the following steps:</p> <ul style="list-style-type: none"> ■ Configure each SCSI device to have a unique SCSI ID, then using the adapter's Advanced Configuration Options setup menu, set the Plug N Play SCAM Support option to Disabled. ■ If there is more than one controller (or an embedded controller), try to use one IRQ per controller.

adp(7D)

- Enable bus mastering for the slots with your host bus adapters, when the choice is given.
- For older disk drives, tape drives, and most CD-ROM devices, make sure the maximum SCSI data transfer speed is set to 5.0 Mbps.
- Enable support for disks larger than 1 Gbyte if applicable.

FILES /kernel/drv/adp.conf Configuration file for the adp driver; there are no user-configurable options in this file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`

Solaris (Intel Platform Edition) Hardware Compatibility List

NOTES Throughout the release, support of additional devices may be added. See the *Solaris (Intel Platform Edition) Hardware Compatibility List* for additional information.

The adp driver supports Logical Unit Number (“LUN”) values of 0 through 15. This range exceeds the standard SCSI-2 requirements which call for support of LUNs 0 through 7.

NAME	adpu320 – Adaptec Ultra320 SCSI host bus adapter driver		
SYNOPSIS	scsi@unit-address		
DESCRIPTION	<p>The adpu320 host bus adapter driver is a SCSI-compliant nexus driver that supports the following Adaptec Ultra320 SCSI Devices:</p> <table border="0"> <tr> <td style="padding-right: 40px;">Chips</td> <td>AIC-7902</td> </tr> </table> <p>The adpu320 driver supports standard functions provided by the SCSI interface, including tagged and untagged queuing, Wide/Fast/Ultra SCSI, and auto request sense. The adpu320 driver does not support linked commands. The adpu320 driver supports hot swap SCSI and hot plug PCI.</p> <p>Additionally, the adpu320 driver supports the following features:</p> <ul style="list-style-type: none"> ■ 64-bit addressing (Dual address Cycle) ■ PCI-X v1.1 operating up to 133MHz and 64bits ■ PCI bus spec v2.2 operating up to 66MHz and 64bits ■ Packetized SCSI at 320 and 160 MB/s ■ QAS ■ DT ■ 40MB/sec in single-ended mode and up to 320MB/sec transfer rate in LVD mode ■ Domain Validation ■ Retained Training Information (RTI) ■ PCI and PCI-X Error handling 	Chips	AIC-7902
Chips	AIC-7902		
Driver Configuration	<p>The adpu320 host bus adapter driver is configured by defining the properties found in adpu320.conf. Properties in the adpu320.conf file that can be modified by the user include: ADPU320_SCSI_RD_STRM, ADPU320_SCSI_NLUN_SUPPORT.</p> <pre> ----- Option: ADPU320_SCSI_RD_STRM=[value] Definition: Enables/disables read streaming negotiation for all drives. Possible Values: 0 (off), 1 (on) Default Value: 0 (off) Option: ADPU320_SCSI_NLUN_SUPPORT=[value] Definition: Enables the number of logical units to be scanned per drive. Possible Values: 1-64 Default Value: 64 ----- </pre> <p><i>If you alter or add driver parameters incorrectly, you can render your system inoperable. Use driver parameters with caution.</i></p>		
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:		

adpu320(7D)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	x86

FILES

/kernel/drv/adpu320
driver module

/kernel/drv/amd64/adpu320
64-bit driver module

/kernel/drv/adpu320.conf
configuration file

/boot/solaris/drivers/notisa.010/adpu320.bef
realmode driver

SEE ALSO

cfgadm(1M), prtconf(1M), attributes(5), scsi_abort(9F),
scsi_hba_attach(9F), scsi_ifgetcap(9F), scsi_reset(9F),
scsi_sync_pkt(9F), scsi_transport(9F), scsi_device(9S),
scsi_extended_sense(9S), scsi_inquiry(9S), scsi_pkt(9S)

Writing Device Drivers

Small Computer System Interface-3 (SCSI-3)

NAME	afb – Elite3D graphics accelerator driver
DESCRIPTION	The afb driver is the device driver for the Sun Elite3D graphics accelerators. The afbdaemonprocess loads the afb microcode at system startup time and during the resume sequence of a suspend-resume cycle.
FILES	<code>/dev/fbs/afb<i>n</i></code> Device special file <code>/usr/lib/afb.unicode</code> afb microcode <code>/usr/sbin/afbdaemon</code> afb microcode loader
SEE ALSO	afbconfig(1M)

amr(7D)

NAME amr – SCSI HBA driver for LSI MegaRAID Controller

DESCRIPTION The amr plain SCSI host bus adapter driver is a SCSI-compliant nexus driver that supports the Dell PERC 3DC/4SC/4DC/4Di and LSI MegaRAID 320-2x SCSI RAID devices.

The amr driver ports from FreeBSD and only supports basic RAID disk I/O functions.

DRIVER CONFIGURATION There are no user configurable parameters available. Please configure your hardware through BIOS.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	x86

FILES /kernel/drv/amr
32-bit ELF kernel module.

/kernel/drv/amd64/amr
kernel module for 64-bit x86 platform

/kernel/drv/amr.conf
driver configuration file (contains no user-configurable options)

SEE ALSO `prtconf(1M)`, `attributes(5)`, `scsi_hba_attach(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_inquiry(9S)`, `scsi_device(9S)`, `scsi_pkt(9S)`

Small Computer System Interface-2 (SCSI-2)

NAME	arp, ARP – Address Resolution Protocol
SYNOPSIS	<pre>#include <sys/fcntl.h> #include <sys/socket.h> #include <net/if_arp.h> #include <netinet/in.h> s = socket(AF_INET, SOCK_DGRAM, 0); d = open ("/dev/arp", oflag);</pre>
DESCRIPTION	<p>ARP is a protocol used to map dynamically between Internet Protocol (IP) and 10Mb/s Ethernet addresses. It is used by all the 10Mb/s Ethernet datalink providers (interface drivers) and it can be used by other datalink providers that support broadcast, such as FDDI and Token Ring. The only network layer supported in this implementation is the Internet Protocol, although ARP is not specific to that protocol.</p> <p>ARP caches IP-to-link-layer address mappings. When an interface requests a mapping for an address not in the cache, ARP queues the message that requires the mapping and broadcasts a message on the associated network requesting the address mapping. If a response is provided, ARP caches the new mapping and transmits any pending message. ARP will queue a maximum of four packets while awaiting a response to a mapping request; it keeps only the four most recently transmitted packets.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The STREAMS device <code>/dev/arp</code> is not a Transport Level Interface (TLI) transport provider and may not be used with the TLI interface.</p> <p>To facilitate communications with systems that do not use ARP, <code>ioctl()</code> requests are provided to enter and delete entries in the IP-to-link address tables.</p> <pre>#include <sys/sockio.h> #include <sys/socket.h> #include <net/if.h> #include <net/if_arp.h> struct arpreq arpreq; ioctl(s, SIOCSARP, (caddr_t)&arpreq); ioctl(s, SIOCGARP, (caddr_t)&arpreq); ioctl(s, SIOCDEARP, (caddr_t)&arpreq);</pre> <p>SIOCSARP, SIOCGARP and SIOCDEARP are BSD compatible <code>ioctl</code>s. These <code>ioctl</code>s do not communicate the mac address length between the user and the kernel (and thus only work for 6 byte wide Ethernet addresses). To manage the ARP cache for media that has different sized mac addresses, use SIOCSXARP, SIOCGXARP and SIOCDXARP <code>ioctl</code>s.</p> <pre>#include <sys/sockio.h> #include <sys/socket.h> #include <net/if.h> #include <net/if_dl.h> #include <net/if_arp.h> struct xarpreq xarpreq;</pre>

arp(7P)

```
ioctl(s, SIOCSXARP, (caddr_t)&xarpreq);
ioctl(s, SIOCGXARP, (caddr_t)&xarpreq);
ioctl(s, SIOCDXARP, (caddr_t)&xarpreq);
```

Each `ioctl()` request takes the same structure as an argument. `SIOCS[X] ARP` sets an ARP entry, `SIOCG[X] ARP` gets an ARP entry, and `SIOC[X] ARP` deletes an ARP entry. These `ioctl()` requests may be applied to any Internet family socket descriptor `s`, or to a descriptor for the ARP device, but only by the privileged user.

The `arpreq` structure contains

```
/*
 * ARP ioctl request
 */
struct arpreq {
    struct sockaddr arp_pa; /* protocol address */
    struct sockaddr arp_ha; /* hardware address */
    int arp_flags; /* flags */
};
```

The `xarpreq` structure contains:

```
/*
 * Extended ARP ioctl request
 */
struct xarpreq {
    struct sockaddr_storage xarp_pa; /* protocol address */
    struct sockaddr_dl xarp_ha; /* hardware address */
    int xarp_flags; /* flags */
};

/* arp_flags field values */
#define ATF_COM 0x2 /* completed entry (arp_ha valid) */
#define ATF_PERM 0x4 /* permanent entry */
#define ATF_PUBL 0x8 /* publish (respond for other host) */
#define ATF_USETRAILERS 0x10 /* send trailer packets to host */
```

The address family for the `[x]arp_pa` `sockaddr` must be `AF_INET`. The only flag bits (`[x]arp_flags`) that may be written are `ATF_PUBL` and `ATF_USETRAILERS`. `ATF_PERM` makes the entry permanent if the `ioctl()` request succeeds. The peculiar nature of the ARP tables may cause the `ioctl()` request to fail if too many permanent IP addresses hash to the same slot. `ATF_PUBL` specifies that the ARP code should respond to ARP requests for the indicated host coming from other machines. This allows a host to act as an “ARP server”, which may be useful in convincing an ARP-only machine to talk to a non-ARP machine.

The address family for the `arp_ha` `sockaddr` must be `AF_UNSPEC`.

Before invoking any of the `SIOC*XARP` `ioctl`s, user code must fill in the `xarp_pa` field with the protocol (IP) address information, similar to the BSD variant. The `SIOC*XARP` `ioctl`s come in two (legal) varieties, depending on `xarp_ha.sdl_nlen`:

1. if `sdl_nlen = 0`, it behaves as an extended BSD `ioctl`. The kernel uses the IP address to determine the network interface.
2. if `(sdl_nlen > 0)` and `(sdl_nlen < LIFNAMSIZ)`, the kernel uses the interface name in `sdl_data[0]` to determine the network interface; `sdl_nlen` represents the length of the string (excluding terminating null character).
3. if `(sdl_nlen >= LIFNAMSIZ)`, an error (`EINVAL`) is flagged from the `ioctl`.

Other than the above, the `xarp_ha` structure should be 0-filled except for `SIOCSXARP`, where the `sdl_alen` field must be set to the size of hardware address length and the hardware address itself must be placed in the `LLADDR/sdl_data[]` area. (`EINVAL` will be returned if user specified `sdl_alen` does not match the address length of the identified interface).

On return from the kernel on a `SIOCGXARP` `ioctl`, the kernel fills in the name of the interface (excluding terminating `NULL`) and its hardware address, one after another, in the `sdl_data/LLADDR` area; if the two are larger than can be held in the 244 byte `sdl_data[]` area, an `ENOSPC` error is returned. Assuming it fits, the kernel will also set `sdl_alen` with the length of hardware address, `sdl_nlen` with the length of name of the interface (excluding terminating `NULL`), `sdl_type` with an `IFT_*` value to indicate the type of the media, `sdl_slen` with 0, `sdl_family` with `AF_LINK` and `sdl_index` (which if not 0) with system given index for the interface. The information returned is very similar to that returned via routing sockets on an `RTM_IFINFO` message.

ARP is also used to negotiate the use of trailer IP encapsulations. Trailers are an alternate encapsulation used to allow efficient packet alignment for large packets despite variable-sized headers. Hosts that wish to receive trailer encapsulations so indicate by sending gratuitous ARP translation replies along with replies to IP requests; trailer encapsulations are also sent in reply to IP translation replies. The negotiation is thus fully symmetrical, in that either host or both may request trailers. The `ATF_USETRAILERS` flag records the receipt of such a reply and enables the transmission of trailer packets to that host.

ARP watches passively for hosts impersonating the local host (that is, a host which responds to an ARP mapping request for the local host's address).

SEE ALSO `arp(1M)`, `ifconfig(1M)`, `if_tcp(7P)`, `inet(7P)`

Leffler, Sam, and Michael Karels, *Trailer Encapsulations*, RFC 893, Network Information Center, SRI International, Menlo Park, California, April 1984.

Plummer, Dave, *An Ethernet Address Resolution Protocol -or- Converting Network Protocol Addresses to 48.bit Ethernet Addresses for Transmission on Ethernet Hardware*, RFC 826, Network Information Center, SRI International, Menlo Park, California, November 1982.

DIAGNOSTICS IP: Hardware address '%x:%x ... %x:%x'

trying to be our address '%d.%d.%d.%d'!

Duplicate IP address. ARP has discovered another host on the local network which responds to mapping requests for the Internet address of this system.

arp(7P)

IP: Proxy ARP problem? Hardware address '%x:%x ... %x:%x'
thinks it is '%d.%d.%d.%d'

This message will appear if `arp(1M)` has been used to create a published entry, and some other host on the local network responds to mapping requests for the published ARP entry.

In the above examples, the "Hardware address" strings include colon (:) separated ascii representations of the link layer addresses, whose lengths depend on the underlying media (for example, 6 bytes for Ethernet).

NAME	asy – asynchronous serial port driver
SYNOPSIS	<pre>#include <fcntl.h> #include <sys/termios.h> open("/dev/term/n", mode); open("/dev/tty/n", mode); open("/dev/cua/n", mode);</pre>
DESCRIPTION	<p>The <code>asy</code> module is a loadable STREAMS driver that provides basic support for Intel-8250, National Semiconductor-16450, 16550, and some 16650 and 16750 and equivalent UARTs connected via the ISA-bus, in addition to basic asynchronous communication support. The <code>asy</code> module supports those <code>termio(7I)</code> device control functions specified by flags in the <code>c_cflag</code> word of the <code>termios</code> structure, and by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code>, <code>INPCK</code>, <code>IXON</code>, <code>IXANY</code>, or <code>IXOFF</code> flags in the <code>c_iflag</code> word of the <code>termios</code> structure. All other <code>termio(7I)</code> functions must be performed by STREAMS modules pushed atop the driver. When a device is opened, the <code>ldterm(7M)</code> and <code>ttcompat(7M)</code> STREAMS modules are automatically pushed on top of the stream, providing the standard <code>termio(7I)</code> interface.</p> <p>The character-special devices <code>/dev/term/a</code>, <code>/dev/term/b</code>, <code>/dev/term/c</code> and <code>/dev/term/d</code> are used to access the four standard serial ports (COM1, COM2, COM3 and COM4 at I/O addresses 3f8, 2f8, 3e8 and 2e8 respectively). Serial ports on non-standard ISA-bus I/O addresses are accessed via the character-special devices <code>/dev/term/0</code>, <code>/dev/term/1</code>, etc. Device names are typically used to provide a logical access point for a dial-in line that is used with a modem.</p> <p>To allow a single tty line to be connected to a modem and used for incoming and outgoing calls, a special feature is available that is controlled by the minor device number. By accessing character-special devices with names of the form <code>/dev/cua/n</code>, it is possible to open a port without the Carrier Detect signal being asserted, either through hardware or an equivalent software mechanism. These devices are commonly known as dial-out lines.</p> <p>Note – This module is affected by the setting of certain eeprom variables, <code>ttya-ignore-cd</code> and <code>ttya-rts-dtr-off</code> (and similarly for <code>ttyb-</code>, <code>ttyc-</code>, and <code>ttyd-</code> parameters). For information on these parameters, see the <code>eeprom(1M)</code> man page.</p> <p>Note – For serial ports on the standard COM1 to COM4 I/O addresses above, the default setting for <code>ttya-ignore-cd</code> and <code>ttya-rts-dtr-off</code> is true. If any of these ports are connected to a modem, these settings should be changed to false. For serial ports on non-standard I/O addresses, the default setting for <code>ttya-ignore-cd</code> and <code>ttya-rts-dtr-off</code> is false.</p>
APPLICATION PROGRAMMING INTERFACE	<p>Once a <code>/dev/cua/n</code> line is opened, the corresponding tty line cannot be opened until the <code>/dev/cua/n</code> line is closed. A blocking open will wait until the <code>/dev/cua/n</code> line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again. A non-blocking open will return an error. If the <code>/dev/ttydn</code> line has been opened successfully (usually only</p>

asy(7D)

when carrier is recognized on the modem), the corresponding `/dev/cua/n` line cannot be opened. This allows a modem to be attached to `/dev/term/[n]` (renamed from `/dev/tty[n]`) and used for dial-in (by enabling the line for login in `/etc/inittab`) or dial-out (by `tip(1)` or `uucp(1C)`) as `/dev/cua/n` when no one is logged in on the line.

IOCTLS

The standard set of `termio ioctl()` calls are supported by `asy`.

Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls.

The input and output line speeds may be set to any speed that is supported by `termio`. The speeds cannot be set independently; for example, when the output speed is set, the input speed is automatically set to the same speed.

When the `asy` module is used to service the serial console port, it supports a `BREAK` condition that allows the system to enter the debugger or the monitor. The `BREAK` condition is generated by hardware and it is usually enabled by default.

A `BREAK` condition originating from erroneous electrical signals cannot be distinguished from one deliberately sent by remote DCE. The Alternate Break sequence can be used as a remedy against this. Due to a risk of incorrect sequence interpretation, `SLIP` and certain other binary protocols should not be run over the serial console port when Alternate Break sequence is in effect. Although `PPP` is a binary protocol, it is able to avoid these sequences using the `ACCM` feature in *RFC 1662*. For Solaris `PPP 4.0`, you do this by adding the following line to the `/etc/ppp/options` file (or other configuration files used for the connection; see `pppd(1M)` for details):

```
asynctmap 0x00002000
```

By default, the Alternate Break sequence is a three character sequence: carriage return, tilde and control-B (`CR ~ CTRL-B`), but may be changed by the driver. For more information on breaking (entering the debugger or monitor), see `kbd(1)` and `kb(7M)`.

ERRORS

An `open()` will fail under the following conditions:

- | | |
|--------------------|---|
| <code>ENXIO</code> | The unit being opened does not exist. |
| <code>EBUSY</code> | The dial-out device is being opened while the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open. |
| <code>EBUSY</code> | The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call. |
| <code>EINTR</code> | The open was interrupted by the delivery of a signal. |

FILES

`/dev/term/[a-d]`
`/dev/term/[012...]`
dial-in tty lines

`/dev/cua/[a-d]`
`/dev/cua/[012...]`
dial-out tty lines

/kernel/drv/amd64/asy
64-bit kernel module for 64-bit x86 platform

/kernel/drv/asy.conf
asy configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `tip(1)`, `kbd(1)`, `uucp(1C)`, `eeprom(1M)`, `pppd(1M)`, `ioctl(2)`, `open(2)`, `termios(3C)`, `attributes(5)`, `ldterm(7M)`, `ttcompat(7M)`, `kb(7M)`, `termio(7I)`

DIAGNOSTICS

<code>asyn</code> : silo overflow.	The hardware overrun occurred before the input character could be serviced.
<code>asyn</code> : ring buffer overflow.	The driver's character input ring buffer overflowed before it could be serviced.

ata(7D)

NAME	ata – AT attachment disk driver
SYNOPSIS	<code>ide@unit-address</code>
DESCRIPTION	<p>The <code>ata</code> driver supports disk and CD-ROM interfaces conforming to the AT Attachment specification including IDE interfaces. Support is provided for both parallel ATA (PATA) and serial ATA (SATA) interfaces. SATA support is limited to version 1 (ATA emulation mode).</p> <p>Refer to the <i>Solaris x86 Hardware Compatibility List</i> for a list of supported controllers.</p>
PRECONFIGURE	<p>A PCI IDE controller can operate in compatibility mode or in PCI-native mode. If more than one controller is present in the system, only one of them can operate in compatibility mode.</p> <p>If two PATA drives share the same controller, you must set one to master and the other to slave. If both a PATA disk drive and a PATA CD-ROM drive utilize the same controller, you can designate the disk drive as the master with the CD-ROM drive as the slave, although this is not mandatory.</p>
Supported Settings	<p>Supported settings for the primary controller when in compatibility mode are:</p> <ul style="list-style-type: none">■ IRQ Level: 14■ I/O Address: 0x1F0 <p>Supported settings for the secondary controller when in compatibility mode are:</p> <ul style="list-style-type: none">■ IRQ Level: 15■ I/O Address: 0x170 <p>Note – When in PCI-native mode, the IRQ and I/O address resources are configured by the system BIOS.</p>
Known Problems and Limitations	<ul style="list-style-type: none">■ This driver does not support any RAID features present on PATA/SATA controller. As a result, you should configure BIOS to select IDE mode rather than RAID mode. Some systems may require updating BIOS to allow switching modes.■ On some systems, the SATA controller needs to have option ROM enabled or BIOS will not consider SATA drives as bootable devices.■ Panasonic LK-MC579B and the Mitsumi FX34005 IDE CD-ROM drives are not supported and cannot be used to install the Solaris operating environment.■ CMD-604 is unable to handle simultaneous I/O on both IDE interfaces. This defect causes the Solaris software to hang if both interfaces are used. Use only the primary IDE interface at address 0x1F0.■ The Solaris Volume Management software does not work with the Sony CDU-55E CD-ROM drive whether configured as device 0 or 1 (master or slave). Comment out the following line in the <code>/etc/vold.conf</code> file to prevent <code>vold</code> from hanging the controller: <pre># use cdrom drive /dev/rdisk/c*s2 dev_cdrom.so cdrom%d</pre>

Direct Memory Access (DMA) and PCI-IDE Systems

- NEC CDR-260/CDR-260R/CDR-273 and Sony CDU-55E ATAPI CD-ROM drives might fail during installation.
- Sony CDU-701 CD-ROM drives must be upgraded to use firmware version 1.0r or later to support booting from the CD.

Direct Memory Access is enabled by default. To disable DMA for the ata driver, do the following steps after you have installed the Solaris operating environment:

1. Run the Solaris (x86 Edition) Device Configuration Assistant from the boot diskette or the installation CD (if your system supports CD-ROM booting). If you are booting from CD, you can interrupt the autoboot process and launch the Device Configuration Assistant by pressing ESC.

Note – After you boot using the boot diskette, the new `ata-dma-enabled` property value is preserved on the diskette. This means that the changed value is in effect each time you use the boot diskette.

2. Press `F2_Continue` to scan for devices.
3. Press `F2_Continue` to display a list of boot devices on the Boot Solaris menu.
4. Go to the View/Edit Property Settings menu.
5. Press `F4_Boot Tasks`, select View/Edit Property Settings, and press `F2_Continue`.
6. Change the value of the `ata-dma-enabled` property. A value of 1 indicates that DMA is enabled and 0 indicates that DMA is disabled.
 - Select the `ata-dma-disabled` property from the list and press `F3_Change`.
 - Type 0 and press `F2_Continue` to disable DMA.
 - Press `F2_Back` to return to the Boot Tasks menu.
 - Press `F3_Back` to return to the Boot Solaris menu.
 - Select the device from which you want to install (network adapter or CD-ROM drive) and press `F2_Continue`.

CONFIGURATION

The driver initializes itself in accordance with the information found in the `ata.conf` configuration file (see below). The following user configurable items are found in the `ata.conf` file:

```
drive0_block_factor
drive1_block_factor
```

ATA controllers support some amount of buffering (blocking). The purpose is to interrupt the host when an entire buffer full of data has been read or written instead of using an interrupt for each sector. This reduces interrupt overhead and significantly increases throughput. The driver interrogates the controller to find the buffer size. Some controllers hang when buffering is used, so the values in the configuration file are used by the driver to reduce the effect of buffering (blocking). The values presented may be chosen from 0x1, 0x2, 0x4, 0x8 and 0x10.

ata(7D)

The values as shipped are set to 0x1, and they can be tuned to increase performance.

If your controller hangs when attempting to use higher block factors, you may be unable to reboot the system. For x86 based systems, it is recommended that the tuning be carried out using a duplicate of the `/platform/i86pc/kernel` directory subtree. This will ensure that a bootable kernel subtree exists in the event of a failed test.

```
ata-revert-to-defaults
revert-<diskmodel>
```

When rebooting or shutting down, the driver can set a feature which allows the drive to return to the power-on settings when the drive receives a software reset (SRST) sequence. If this property is present and set to 1, the driver will set the feature to revert to defaults during reset. Setting this property to 1 may prevent some systems from soft-rebooting and would require cycling the power to boot the system. If this property is not present the system will not set the feature to revert to defaults during reset.

To determine the string to substitute for `<diskmodel>`, boot your system (you may have to press the reset button or power-cycle) and then view `/var/adm/messages`. Look for the string "IDE device at targ" or "ATAPI device at targ." The next line will contain the word "model" followed by the model number and a comma. Ignore all characters except letters, digits, ".", "_", and "-". Change uppercase letters to lower case. If the string `revert-<diskmodel>` is longer than 31 characters, use only the first 31 characters.

EXAMPLES

EXAMPLE 1 Sample ata Configuration File

```
# for higher performance - set block factor to 16
drive0_block_factor=0x1 drive1_block_factor=0x1
max_transfer=0x100
flow_control="dmult" queue="qsort" disk="dadk" ;
```

EXAMPLE 2 Revert to defaults property

```
revert-st320420a=1;
```

Output of `/var/adm/messages`:

```
Aug 17 06:49:43 caesar ata:[ID 640982 kern.info] IDE device at targ 0,
                    lun 0 lastlun 0x0
Aug 17 06:49:43 caesar ata:[ID 521533 kern.info] model ST320420A, stat
```

FILES

`/platform/i86pc/kernel/drv/ata`
Device file for 32-bit x86 platforms.

`/platform/i86pc/kernel/drv/amd64/ata`
Device file for 64-bit x86 platforms.

`/platform/i86pc/kernel/drv/ata.conf`
Configuration file for 32-bit x86 platforms.

`/platform/i86pc/kernel/drv/amd64/ata.conf`
Configuration file for 64-bit x86 platforms.

/boot/solaris/bootenv.rc
Boot environment variables file for Solaris x86.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`

INCITS T13 ATA/ATAPI-7 specifications

audio(7I)

NAME	audio – generic audio device interface
SYNOPSIS	<pre>#include <sys/audio.h></pre>
OVERVIEW	<p>An audio device is used to play and/or record a stream of audio data. Since a specific audio device may not support all functionality described below, refer to the device-specific manual pages for a complete description of each hardware device. An application can use the <code>AUDIO_GETDEV</code> <code>ioctl(2)</code> to determine the current audio hardware associated with <code>/dev/audio</code>.</p>
AUDIO FORMATS	<p>Digital audio data represents a quantized approximation of an analog audio signal waveform. In the simplest case, these quantized numbers represent the amplitude of the input waveform at particular sampling intervals. To achieve the best approximation of an input signal, the highest possible sampling frequency and precision should be used. However, increased accuracy comes at a cost of increased data storage requirements. For instance, one minute of monaural audio recorded in μ-Law format (pronounced <i>mew-law</i>) at 8 KHz requires nearly 0.5 megabytes of storage, while the standard Compact Disc audio format (stereo 16-bit linear PCM data sampled at 44.1 KHz) requires approximately 10 megabytes per minute.</p> <p>Audio data may be represented in several different formats. An audio device's current audio data format can be determined by using the <code>AUDIO_GETINFO</code> <code>ioctl(2)</code> described below.</p> <p>An audio data format is characterized in the audio driver by four parameters: Sample Rate, Encoding, Precision, and Channels. Refer to the device-specific manual pages for a list of the audio formats that each device supports. In addition to the formats that the audio device supports directly, other formats provide higher data compression. Applications may convert audio data to and from these formats when playing or recording.</p>
Sample Rate	<p>Sample rate is a number that represents the sampling frequency (in samples per second) of the audio data.</p>
Encodings	<p>An encoding parameter specifies the audio data representation. μ-Law encoding corresponds to <i>CCITT G.711</i>, and is the standard for voice data used by telephone companies in the United States, Canada, and Japan. A-Law encoding is also part of <i>CCITT G.711</i> and is the standard encoding for telephony elsewhere in the world. A-Law and μ-Law audio data are sampled at a rate of 8000 samples per second with 12-bit precision, with the data compressed to 8-bit samples. The resulting audio data quality is equivalent to that of standard analog telephone service.</p> <p>Linear Pulse Code Modulation (PCM) is an uncompressed, signed audio format in which sample values are directly proportional to audio signal voltages. Each sample is a 2's complement number that represents a positive or negative amplitude.</p>
Precision	<p>Precision indicates the number of bits used to store each audio sample. For instance, μ-Law and A-Law data are stored with 8-bit precision. PCM data may be stored at various precisions, though 16-bit is the most common.</p>

Channels	Multiple channels of audio may be interleaved at sample boundaries. A sample frame consists of a single sample from each active channel. For example, a sample frame of stereo 16-bit PCM data consists of 2 16-bit samples, corresponding to the left and right channel data.
DESCRIPTION	<p>The device <code>/dev/audio</code> is a device driver that dispatches audio requests to the appropriate underlying audio hardware. The audio driver is implemented as a STREAMS driver. In order to record audio input, applications <code>open(2)</code> the <code>/dev/audio</code> device and read data from it using the <code>read(2)</code> system call. Similarly, sound data is queued to the audio output port by using the <code>write(2)</code> system call. Device configuration is performed using the <code>ioctl(2)</code> interface.</p> <p>Alternatively, opening <code>/dev/audio</code> may open a mixing audio driver that provides a super set of this audio interface. The audio mixer removes the exclusive resource restriction, allowing multiple processes to play and record audio at the same time. See the mixer(7I) and audio_support(7I) manual pages for more information.</p> <p>Because some systems may contain more than one audio device, application writers are encouraged to query the <code>AUDIODEV</code> environment variable. If this variable is present in the environment, its value should identify the path name of the default audio device.</p>
Opening the Audio Device	<p>The audio device is treated as an exclusive resource, meaning that only one process can open the device at a time. However, if the <code>DUPLEX</code> bit is set in the <code>hw_features</code> field of the audio information structure, two processes may simultaneously access the device. This allows one process to open the device as read-only and a second process to open it as write-only. See below for details.</p> <p>When a process cannot open <code>/dev/audio</code> because the device is busy:</p> <ul style="list-style-type: none"> ■ if either the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flags are set in the <code>open()</code> <i>oflag</i> argument, then <code>-1</code> is immediately returned, with <i>errno</i> set to <code>EBUSY</code>. ■ if neither the <code>O_NDELAY</code> nor the <code>O_NONBLOCK</code> flag are set, then <code>open()</code> hangs until the device is available or a signal is delivered to the process, in which case a <code>-1</code> is returned with <i>errno</i> set to <code>EINTR</code>. This allows a process to block in the <code>open</code> call while waiting for the audio device to become available. <p>Upon the initial <code>open()</code> of the audio device, the driver resets the data format of the device to the default state of 8-bit, 8Khz, mono u-Law data. If the device is already open and a different audio format is set, this will not be possible on some devices. With the exception of some devices that only support a limited number of sample rates, audio applications should explicitly set the encoding characteristics to match the audio data requirements rather than depend on the default configuration.</p> <p>Since the audio device grants exclusive read or write access to a single process at a time, long-lived audio applications may choose to close the device when they enter an idle state and reopen it when required. The <i>play.waiting</i> and <i>record.waiting</i> flags in the audio information structure (see below) provide an indication that another process has requested access to the device. For instance, a background audio output process may choose to relinquish the audio device whenever another process requests write access.</p>

audio(7I)

Recording Audio Data

The `read()` system call copies data from the system's buffers to the application. Ordinarily, `read()` blocks until the user buffer is filled. The `I_NREAD` ioctl (see [streamio\(7I\)](#)) may be used to determine the amount of data that may be read without blocking. The device may alternatively be set to a non-blocking mode, in which case `read()` completes immediately, but may return fewer bytes than requested. Refer to the `read(2)` manual page for a complete description of this behavior.

When the audio device is opened with read access, the device driver immediately starts buffering audio input data. Since this consumes system resources, processes that do not record audio data should open the device write-only (`O_WRONLY`).

The transfer of input data to STREAMS buffers may be paused (or resumed) by using the `AUDIO_SETINFO` ioctl to set (or clear) the `record.pause` flag in the audio information structure (see below). All unread input data in the STREAMS queue may be discarded by using the `I_FLUSH_STREAMS` ioctl (see [streamio\(7I\)](#)). When changing record parameters, the input stream should be paused and flushed before the change, and resumed afterward. Otherwise, subsequent reads may return samples in the old format followed by samples in the new format. This is particularly important when new parameters result in a changed sample size.

Input data can accumulate in STREAMS buffers very quickly. At a minimum, it will accumulate at 8000 bytes per second for 8-bit, 8 KHz, mono, u-Law data. If the device is configured for 16-bit linear or higher sample rates, it will accumulate even faster. If the application that consumes the data cannot keep up with this data rate, the STREAMS queue may become full. When this occurs, the `record.error` flag is set in the audio information structure and input sampling ceases until there is room in the input queue for additional data. In such cases, the input data stream contains a discontinuity. For this reason, audio recording applications should open the audio device when they are prepared to begin reading data, rather than at the start of extensive initialization.

Playing Audio Data

The `write()` system call copies data from an application's buffer to the STREAMS output queue. Ordinarily, `write()` blocks until the entire user buffer is transferred. The device may alternatively be set to a non-blocking mode, in which case `write()` completes immediately, but may have transferred fewer bytes than requested (see `write(2)`).

Although `write()` returns when the data is successfully queued, the actual completion of audio output may take considerably longer. The `AUDIO_DRAIN` ioctl may be issued to allow an application to block until all of the queued output data has been played. Alternatively, a process may request asynchronous notification of output completion by writing a zero-length buffer (end-of-file record) to the output stream. When such a buffer has been processed, the `play.eof` flag in the audio information structure (see below) is incremented.

The final `close(2)` of the file descriptor hangs until all of the audio output has drained. If a signal interrupts the `close()`, or if the process exits without closing the device, any remaining data queued for audio output is flushed and the device is closed immediately.

	<p>The consumption of output data may be paused (or resumed) by using the <code>AUDIO_SETINFO</code> <code>ioctl</code> to set (or clear) the <code>play.pause</code> flag in the audio information structure. Queued output data may be discarded by using the <code>I_FLUSH STREAMS</code> <code>ioctl</code>. (See <code>streamio(7I)</code>).</p>
	<p>Output data is played from the <code>STREAMS</code> buffers at a default rate of at least 8000 bytes per second for μ-Law, A-Law or 8-bit PCM data (faster for 16-bit linear data or higher sampling rates). If the output queue becomes empty, the <code>play.error</code> flag is set in the audio information structure and output is stopped until additional data is written. If an application attempts to write a number of bytes that is not a multiple of the current sample frame size, an error is generated and the bad data is thrown away. Additional writes are allowed.</p>
Asynchronous I/O	<p>The <code>I_SETSIG STREAMS</code> <code>ioctl</code> enables asynchronous notification, through the <code>SIGPOLL</code> signal, of input and output ready condition changes. The <code>O_NONBLOCK</code> flag may be set using the <code>F_SETFL</code> <code>fcntl(2)</code> to enable non-blocking <code>read()</code> and <code>write()</code> requests. This is normally sufficient for applications to maintain an audio stream in the background.</p>
Audio Control Pseudo-Device	<p>It is sometimes convenient to have an application, such as a volume control panel, modify certain characteristics of the audio device while it is being used by an unrelated process. The <code>/dev/audiocctl</code> pseudo-device is provided for this purpose. Any number of processes may open <code>/dev/audiocctl</code> simultaneously. However, <code>read()</code> and <code>write()</code> system calls are ignored by <code>/dev/audiocctl</code>. The <code>AUDIO_GETINFO</code> and <code>AUDIO_SETINFO</code> <code>ioctl</code> commands may be issued to <code>/dev/audiocctl</code> to determine the status or alter the behavior of <code>/dev/audio</code>. Note: In general, the audio control device name is constructed by appending the letters "ctl" to the path name of the audio device.</p>
Audio Status Change Notification	<p>Applications that open the audio control pseudo-device may request asynchronous notification of changes in the state of the audio device by setting the <code>S_MSG</code> flag in an <code>I_SETSIG STREAMS</code> <code>ioctl</code>. Such processes receive a <code>SIGPOLL</code> signal when any of the following events occur:</p> <ul style="list-style-type: none"> ■ An <code>AUDIO_SETINFO</code> <code>ioctl</code> has altered the device state. ■ An input overflow or output underflow has occurred. ■ An end-of-file record (zero-length buffer) has been processed on output. ■ An <code>open()</code> or <code>close()</code> of <code>/dev/audio</code> has altered the device state. ■ An external event (such as speakerbox's volume control) has altered the device state.
IOCTLS	
Audio Information Structure	<p>The state of the audio device may be polled or modified using the <code>AUDIO_GETINFO</code> and <code>AUDIO_SETINFO</code> <code>ioctl</code> commands. These commands operate on the <code>audio_info</code> structure as defined, in <code><sys/audio.h></code>, as follows:</p>
	<pre>/* * This structure contains state information for audio device</pre>

audio(7I)

```

* IO streams
*/

struct audio_prinfo {
    /*
     * The following values describe the
     * audio data encoding
     */
    uint_t sample_rate; /* samples per second */
    uint_t channels;    /* number of interleaved channels */
    uint_t precision;  /* number of bits per sample */
    uint_t encoding;   /* data encoding method */

    /*
     * The following values control audio device
     * configuration
     */
    uint_t gain;        /* volume level */
    uint_t port;        /* selected I/O port */
    uint_t buffer_size; /* I/O buffer size */

    /*
     * The following values describe the current device
     * state
     */
    uint_t samples;    /* number of samples converted */
    uint_t eof;        /* End Of File counter (play only) */
    uchar_t pause;    /* non-zero if paused, zero to resume */
    uchar_t error;    /* non-zero if overflow/underflow */
    uchar_t waiting;  /* non-zero if a process wants access */
    uchar_t balance;  /* stereo channel balance */

    /*
     * The following values are read-only device state
     * information
     */
    uchar_t open;     /* non-zero if open access granted */
    uchar_t active;   /* non-zero if I/O active */
    uint_t avail_ports; /* available I/O ports */
    uint_t mod_ports; /* modifiable I/O ports */
};
typedef struct audio_prinfo audio_prinfo_t;

/*
 * This structure is used in AUDIO_GETINFO and AUDIO_SETINFO ioctl
 * commands
 */
struct audio_info {
    audio_prinfo_t record; /* input status info */
    audio_prinfo_t play;  /* output status info */
    uint_t monitor_gain; /* input to output mix */
    uchar_t output_muted; /* non-zero if output muted */
    uint_t hw_features; /* supported H/W features */
    uint_t sw_features; /* supported S/W features */
    uint_t sw_features_enabled; /* supported S/W features enabled */
};
typedef struct audio_info audio_info_t;

```

```

/* Audio encoding types */
#define AUDIO_ENCODING_ULAW    (1) /* u-Law encoding */
#define AUDIO_ENCODING_ALAW    (2) /* A-Law encoding */
#define AUDIO_ENCODING_LINEAR  (3) /* Signed Linear PCM encoding */

/*
 * These ranges apply to record, play, and
 * monitor gain values
 */
#define AUDIO_MIN_GAIN (0)      /* minimum gain value */
#define AUDIO_MAX_GAIN (255)   /* maximum gain value */

/*
 * These values apply to the balance field to adjust channel
 * gain values
 */
#define AUDIO_LEFT_BALANCE    (0) /* left channel only */
#define AUDIO_MID_BALANCE    (32) /* equal left/right balance */
#define AUDIO_RIGHT_BALANCE  (64) /* right channel only */

/*
 * Define some convenient audio port names
 * (for port, avail_ports and mod_ports)
 */

/* output ports (several might be enabled at once) */
#define AUDIO_SPEAKER    (0x01) /* built-in speaker */
#define AUDIO_HEADPHONE (0x02) /* headphone jack */
#define AUDIO_LINE_OUT   (0x04) /* line out */
#define AUDIO_SPDIF_OUT  (0x08) /* SPDIF port */
#define AUDIO_AUX1_OUT   (0x10) /* aux1 out */
#define AUDIO_AUX2_OUT   (0x20) /* aux2 out */

/* input ports (usually only one may be
 * enabled at a time)
 */
#define AUDIO_MICROPHONE (0x01) /* microphone */
#define AUDIO_LINE_IN    (0x02) /* line in */
#define AUDIO_CD         (0x04) /* on-board CD inputs */
#define AUDIO_SPDIF_IN   (0x08) /* SPDIF port */
#define AUDIO_AUX1_IN    (0x10) /* aux1 in */
#define AUDIO_AUX2_IN    (0x20) /* aux2 in */
#define AUDIO_CODEC_LOOPB_IN (0x40) /* Codec inter.loopback */

/* These defines are for hardware features */
#define AUDIO_HWFEATURE_DUPLEX (0x00000001u)
    /*simult. play & cap. supported */

#define AUDIO_HWFEATURE_MSCODEC (0x00000002u)
    /* multi-stream Codec */

/* These defines are for software features */
#define AUDIO_SWFEATURE_MIXER (0x00000001u)
    /* audio mixer audio pers. mod. */

/*
 * Parameter for the AUDIO_GETDEV ioctl
 * to determine current audio devices

```

audio(7I)

```
*/  
  
#define MAX_AUDIO_DEV_LEN      (16)  
struct audio_device {  
    char name[MAX_AUDIO_DEV_LEN];  
    char version[MAX_AUDIO_DEV_LEN];  
    char config[MAX_AUDIO_DEV_LEN];  
};  
typedef struct audio_device audio_device_t;
```

The *play.gain* and *record.gain* fields specify the output and input volume levels. A value of `AUDIO_MAX_GAIN` indicates maximum volume. Audio output may also be temporarily muted by setting a non-zero value in the *output_muted* field. Clearing this field restores audio output to the normal state. Most audio devices allow input data to be monitored by mixing audio input onto the output channel. The *monitor_gain* field controls the level of this feedback path.

The *play.port* field controls the output path for the audio device. It can be set to either `AUDIO_SPEAKER` (built-in speaker), `AUDIO_HEADPHONE` (headphone jack), `AUDIO_LINE_OUT` (line-out port), `AUDIO_AUX1_OUT` (auxiliary1 out), or `AUDIO_AUX2_OUT` (auxiliary2 out). For some devices, it may be set to a combination of these ports. The *play.avail_ports* field returns the set of output ports that are currently accessible. The *play.mod_ports* field returns the set of output ports that may be turned on and off. If a port is missing from *play.mod_ports* then that port is assumed to always be on.

The *record.port* field controls the input path for the audio device. It can be either `AUDIO_MICROPHONE` (microphone jack), `AUDIO_LINE_IN` (line-out port), `AUDIO_CD` (internal CD-ROM), `AUDIO_AUX1_IN` (auxiliary1 in), `AUDIO_AUX2_IN` (auxiliary2 in), or `AUDIO_CODEC_LOOPBACK_IN` (internal loopback). The *record.avail_ports* field returns the set of input ports that are currently accessible. The *record.mod_ports* field returns the set of input ports that may be turned on and off. If a port is missing from *record.mod_ports*, it is assumed to always be on. Input ports are considered to be mutually exclusive.

The *play.balance* and *record.balance* fields are used to control the volume between the left and right channels when manipulating stereo data. When the value is set between `AUDIO_LEFT_BALANCE` and `AUDIO_MID_BALANCE`, the right channel volume will be reduced in proportion to the *balance* value. Conversely, when *balance* is set between `AUDIO_MID_BALANCE` and `AUDIO_RIGHT_BALANCE`, the left channel will be proportionally reduced.

The *play.pause* and *record.pause* flags may be used to pause and resume the transfer of data between the audio device and the STREAMS buffers. The *play.error* and *record.error* flags indicate that data underflow or overflow has occurred. The *play.active* and *record.active* flags indicate that data transfer is currently active in the corresponding direction.

The *play.open* and *record.open* flags indicate that the device is currently open with the corresponding access permission. The *play.waiting* and *record.waiting* flags provide an indication that a process may be waiting to access the device. These flags are set automatically when a process blocks on `open()`, though they may also be set using the `AUDIO_SETINFO` `ioctl` command. They are cleared only when a process relinquishes access by closing the device.

The *play.samples* and *record.samples* fields are zeroed at `open()` and are incremented each time a data sample is copied to or from the associated STREAMS queue. Some audio drivers may be limited to counting buffers of samples, instead of single samples for their *samples* accounting. For this reason, applications should not assume that the *samples* fields contain a perfectly accurate count. The *play.eof* field increments whenever a zero-length output buffer is synchronously processed. Applications may use this field to detect the completion of particular segments of audio output.

The *record.buffer_size* field controls the amount of input data that is buffered in the device driver during record operations. Applications that have particular requirements for low latency should set the value appropriately. Note however that smaller input buffer sizes may result in higher system overhead. The value of this field is specified in bytes and drivers will constrain it to be a multiple of the current sample frame size. Some drivers may place other requirements on the value of this field. Refer to the audio device-specific manual page for more details. If an application changes the format of the audio device and does not modify the *record.buffer_size* field, the device driver may use a default value to compensate for the new data rate. Therefore, if an application is going to modify this field, it should modify it during or after the format change itself, not before. When changing the *record.buffer_size* parameters, the input stream should be paused and flushed before the change, and resumed afterward. Otherwise, subsequent reads may return samples in the old format followed by samples in the new format. This is particularly important when new parameters result in a changed sample size. If you change the *record.buffer_size* for the first packet, this protocol must be followed or the first buffer will be the default buffer size for the device, followed by packets of the requested change size.

The *record.buffer_size* field may be modified only on the `/dev/audio` device by processes that have it opened for reading.

The *play.buffer_size* field is currently not supported.

The audio data format is indicated by the *sample_rate*, *channels*, *precision*, and *encoding* fields. The values of these fields correspond to the descriptions in the AUDIO FORMATS section above. Refer to the audio device-specific manual pages for a list of supported data format combinations.

The data format fields may be modified only on the `/dev/audio` device. Some audio hardware may constrain the input and output data formats to be identical. If this is the case, the data format may not be changed if multiple processes have opened the audio device. As a result, a process should check that the `ioctl()` does not fail when it attempts to set the data format.

audio(7I)

If the parameter changes requested by an `AUDIO_SETINFO` `ioctl` cannot all be accommodated, `ioctl()` will return with `errno` set to `EINVAL` and no changes will be made to the device state.

Streamio IOCTLS

All of the [streamio\(7I\)](#) `ioctl` commands may be issued for the `/dev/audio` device. Because the `/dev/audioctl` device has its own STREAMS queues, most of these commands neither modify nor report the state of `/dev/audio` if issued for the `/dev/audioctl` device. The `I_SETSIG` `ioctl` may be issued for `/dev/audioctl` to enable the notification of audio status changes, as described above.

Audio IOCTLS

The audio device additionally supports the following `ioctl` commands:

- | | |
|----------------------------|---|
| <code>AUDIO_DRAIN</code> | The argument is ignored. This command suspends the calling process until the output STREAMS queue is empty, or until a signal is delivered to the calling process. It may not be issued for the <code>/dev/audioctl</code> device. An implicit <code>AUDIO_DRAIN</code> is performed on the final <code>close()</code> of <code>/dev/audio</code> . |
| <code>AUDIO_GETDEV</code> | The argument is a pointer to an <code>audio_device_t</code> structure. This command may be issued for either <code>/dev/audio</code> or <code>/dev/audioctl</code> . The returned value in the <code>name</code> field will be a string that will identify the current <code>/dev/audio</code> hardware device, the value in <code>version</code> will be a string indicating the current version of the hardware, and <code>config</code> will be a device-specific string identifying the properties of the audio stream associated with that file descriptor. Refer to the audio device-specific manual pages to determine the actual strings returned by the device driver. |
| <code>AUDIO_GETINFO</code> | The argument is a pointer to an <code>audio_info_t</code> structure. This command may be issued for either <code>/dev/audio</code> or <code>/dev/audioctl</code> . The current state of the <code>/dev/audio</code> device is returned in the structure. |
| <code>AUDIO_SETINFO</code> | The argument is a pointer to an <code>audio_info_t</code> structure. This command may be issued for either the <code>/dev/audio</code> or the <code>/dev/audioctl</code> device with some restrictions. This command configures the audio device according to the supplied structure and overwrites the existing structure with the new state of the device. Note: The <code>play.samples</code> , <code>record.samples</code> , <code>play.error</code> , <code>record.error</code> , and <code>play.eof</code> fields are modified to reflect the state of the device when the <code>AUDIO_SETINFO</code> is issued. This allows programs to automatically modify these fields while retrieving the previous value. |

Certain fields in the audio information structure, such as the `pause` flags, are treated as read-only when `/dev/audio` is not open with the corresponding access permission. Other fields, such as the gain levels and encoding information, may have a restricted set of acceptable values. Applications that attempt to modify such fields should check the returned values to be sure that the corresponding change took effect. The `sample_rate`, `channels`, `precision`, and `encoding` fields treated as read-only for

`/dev/audiocctl`, so that applications can be guaranteed that the existing audio format will stay in place until they relinquish the audio device. `AUDIO_SETINFO` will return `EINVAL` when the desired configuration is not possible, or `EBUSY` when another process has control of the audio device.

Once set, the following values persist through subsequent `open()` and `close()` calls of the device and automatic device unloads: *play.gain*, *record.gain*, *play.balance*, *record.balance*, *play.port*, *record.port*, *monitor_gain* and *output_muted*. For the `dbri` driver, an automatic device driver unload resets these parameters to their default values on the next load. All other state is reset when the corresponding I/O stream of `/dev/audio` is closed.

The `audio_info_t` structure may be initialized through the use of the `AUDIO_INITINFO` macro. This macro sets all fields in the structure to values that are ignored by the `AUDIO_SETINFO` command. For instance, the following code switches the output port from the built-in speaker to the headphone jack without modifying any other audio parameters:

```
audio_info_t  info;
AUDIO_INITINFO(&info);
info.play.port = AUDIO_HEADPHONE;
err = ioctl(audio_fd, AUDIO_SETINFO, &info);
```

This technique eliminates problems associated with using a sequence of `AUDIO_GETINFO` followed by `AUDIO_SETINFO`.

ERRORS

An `open()` will fail if:

<code>EBUSY</code>	The requested play or record access is busy and either the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flag was set in the <code>open()</code> request.
<code>EINTR</code>	The requested play or record access is busy and a signal interrupted the <code>open()</code> request.

An `ioctl()` will fail if:

<code>EINVAL</code>	The parameter changes requested in the <code>AUDIO_SETINFO</code> <code>ioctl</code> are invalid or are not supported by the device.
<code>EBUSY</code>	The parameter changes requested in the <code>AUDIO_SETINFO</code> <code>ioctl</code> could not be made because another process has the device open and is using a different format.

FILES

The physical audio device names are system dependent and are rarely used by programmers. Programmers should use the generic device names listed below.

```
/dev/audio
    symbolic link to the system's primary audio device

/dev/audiocctl
    symbolic link to the control device for /dev/audio
```

audio(7I)

/dev/sound/0
first audio device in the system

/dev/sound/0ctl
audio control device for /dev/sound/0

/usr/share/audio/samples
audio files

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86
Availability	SUNWcsu, SUNWcsxu, SUNWaudd, SUNWauddx, SUNWaudh
Stability Level	Evolving

SEE ALSO `close(2)`, `fcntl(2)`, `ioctl(2)`, `open(2)`, `poll(2)`, `read(2)`, `write(2)`, `attributes(5)`, `audiocs(7D)`, `audioens(7D)`, `audiots(7D)`, `dbri(7D)`, `sbpro(7D)`, `usb_ac(7D)`, `audio_support(7I)`, `mixer(7I)`, `streamio(7I)`

BUGS Due to a *feature* of the STREAMS implementation, programs that are terminated or exit without closing the `audio` device may hang for a short period while audio output drains. In general, programs that produce audio output should catch the `SIGINT` signal and flush the output stream before exiting.

On LX machines running Solaris 2.3, catting a demo audio file to the audio device `/dev/audio` does not work. Use the `audioplay` command on LX machines instead of `cat`.

FUTURE DIRECTIONS Future audio drivers should use the `mixer(7I)` audio device to gain access to new features.

NAME	audio810 – Amd 8111 south bridge audio digital controller interface																
DESCRIPTION	The audio810 device uses the Amd 8111 audio digital controller and an AC-97 Codec to implement the audio device interface.																
APPLICATION PROGRAM INTERFACE	This interface is described in the mixer(7I) and audio(7I) man pages.																
Driver Version	Applications that open /dev/audio may use the AUDIO_GETDEV ioctl(2) to determine which audio device is being used. The audio810 driver returns the string SUNW, audio810 in the name field of the audio_device structure. The version field contains a letter "a" and the config field contains the string onboard1.																
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Platform</th> <th>Ver.</th> <th>Out</th> <th>H. Phone</th> <th>Spkr.</th> <th>In</th> <th>Mic.</th> <th>CD</th> </tr> </thead> <tbody> <tr> <td>Sun W1100z/ W2100z workstation</td> <td>a</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> </tr> </tbody> </table>	Platform	Ver.	Out	H. Phone	Spkr.	In	Mic.	CD	Sun W1100z/ W2100z workstation	a	Y	Y	Y	Y	Y	Y
Platform	Ver.	Out	H. Phone	Spkr.	In	Mic.	CD										
Sun W1100z/ W2100z workstation	a	Y	Y	Y	Y	Y	Y										
	Key to the above table: Ver. = Version. Out = Line Out. H. Phone = Head Phone. Spkr = Internal Speaker. In = Line In. Mic. = Microphone. CD = CD-ROM.																
	The audio810 device provides support for the internal speaker, headphone, line out, line in, and microphone. The play.mod_ports and record.mod_ports fields of the audio_info structure (see audio(7I)) indicate which ports may be manipulated.																
Audio Mixer Mode	This driver supports the mixer mode only. (Compat mode is not supported).																
Audio Data Formats	The audio810 device supports 16-bit linear encodings in stereo. With the mixer enabled, a continuous range of sample rates from 5510 to 48000 Hz is supported. With the mixer disabled, only sample rate 48000 Hz is supported. This is limited by hardware capability. The device can be opened for simultaneous play and record whether the mixer is enabled or not.																
Sample Granularity	Because the audio810 device manipulates buffers of audio data, the reported input and output sample counts will vary at any given time from the actual sample count by no more than the size of the buffers the audio810 driver is transferring. In general, programs should not rely on the absolute accuracy of the play.samples and record.samples fields of the audio_info structure.																
Interrupt Rate	The driver determines how often play and record interrupts should take place. For playing audio, this determines how often and how much audio is requested from the audio mixer. The impact of interrupts on recording is minimal. However, if a very small read buffer size is set, the record interrupt rate should be increased to prevent the buffer from overflowing. The play and record interrupt rates are tunable in the /kernel/drv/audio810.conf file.																
Audio Status Change Notification	As described in the audio(7I) and mixer(7I) man pages, it is possible to request asynchronous notification of changes in the state of an audio device.																

audio810(7D)

ERRORS audio810 errors are described in the audio(7I) man page.

FILES

- /dev/audio
Symbolic link to the system's primary audio device. (Not necessarily an audio810 audio device).
- /dev/audioctrl
Control device for the primary audio device.
- /dev/sound/0
Represents the first audio device on the system. (Not necessarily an audio810 audio device).
- /dev/sound/octrl
Audio control for /dev/sound/0.
- /usr/share/audio/samples
Audio sample files.
- /kernel/drv/audio810.conf
audio810 driver configuration file.

ATTRIBUTES See attributes(5) for a descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PC-based system
Availability	SUNWaudd, SUNWauda
Stability level	Evolving

SEE ALSO mixerctl(1), sdtaudiocontrol(1), ioctl(2), attributes(5), audio(7I), mixer(7I), streamio(7I)

AMD-8111 HyperTransport I/O Hub Data Sheet — Advanced Micro Devices Inc.

ALC655 Specification — Realtek Inc.

<http://www.sun.com/io>

DIAGNOSTICS In addition to being logged, the following messages may appear on the system console:

```
init_state() play interrupt rate set too low
```

The play interrupt rate in audio810.conf is set too low. It has been reset to the rate specified in the message. Update audio810.conf to a higher play interrupt rate.

```
init_state() play interrupt rate set too high
```

The play interrupt rate set in audio810.conf is set too high. It has been reset to the rate specified in the message. Update audio810.conf to a lower play interrupt rate.

audio810(7D)

init_state() record interrupt rate set too low

The record interrupt rate in audio810.conf is set too low. It has been reset to the rate specified in the message. Update audio810.conf to a higher record interrupt rate.

init_state() record interrupt rate set too high

The record interrupt rate in audio810.conf is set too high. It has been reset to the rate specified in the message. Update audio810.conf to a lower record interrupt rate.

audiocs(7D)

NAME	audiocs – Crystal Semiconductor 4231 Audio driver																																																
DESCRIPTION	The <code>audiocs</code> driver supports the Crystal Semiconductor 4231 Codec to implement the audio device interface.																																																
APPLICATION PROGRAM INTERFACE	This interface is described in the <code>audio(7I)</code> and <code>mixer(7I)</code> man pages.																																																
Driver Version	Applications that open <code>/dev/audio</code> may use the <code>AUDIO_GETDEV</code> <code>ioctl(2)</code> to determine which audio device is being used. The <code>audiocs</code> driver will return the string <code>SUNW,CS4231</code> in the <code>name</code> field of the <code>audio_device</code> structure. The <code>version</code> field contains a letter (defined in the table below) and the <code>config</code> field contains the string <code>onboard1</code> .																																																
	<table border="1"> <thead> <tr> <th>Platform</th> <th>Ver.</th> <th>Out</th> <th>H. Phone</th> <th>Speaker</th> <th>In</th> <th>Mic</th> <th>CD</th> </tr> </thead> <tbody> <tr> <td>SS-4/5</td> <td>a</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> </tr> <tr> <td>Ultra-2</td> <td>b</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>N</td> </tr> <tr> <td>Ultra-450</td> <td>f</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>N</td> </tr> <tr> <td>Ultra-30/60/80</td> <td>g</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>N</td> </tr> <tr> <td>Ultra-5/10</td> <td>h</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> </tr> </tbody> </table>	Platform	Ver.	Out	H. Phone	Speaker	In	Mic	CD	SS-4/5	a	Y	Y	Y	Y	Y	Y	Ultra-2	b	Y	Y	Y	Y	Y	N	Ultra-450	f	Y	Y	Y	Y	Y	N	Ultra-30/60/80	g	Y	Y	Y	Y	Y	N	Ultra-5/10	h	Y	Y	Y	Y	Y	Y
Platform	Ver.	Out	H. Phone	Speaker	In	Mic	CD																																										
SS-4/5	a	Y	Y	Y	Y	Y	Y																																										
Ultra-2	b	Y	Y	Y	Y	Y	N																																										
Ultra-450	f	Y	Y	Y	Y	Y	N																																										
Ultra-30/60/80	g	Y	Y	Y	Y	Y	N																																										
Ultra-5/10	h	Y	Y	Y	Y	Y	Y																																										
	Key to the above table: Ver = Version. Out = Line Out. Speaker = Internal Speaker. In = Line In. Mic. = Microphone. CD = CD-ROM.																																																
	The <code>audiocs</code> device provides support for the internal speaker, headphone, line out, line in, microphone, and on some platforms, internal CD-ROM audio in. The <code>play.avail_ports</code> and <code>record.avail_ports</code> fields retrieved by the <code>AUDIO_GETINFO</code> <code>ioctl</code> (see <code>audio(7I)</code>) report which ports are available. The <code>play.mod_ports</code> and <code>record.mod_ports</code> fields indicate which ports may be manipulated.																																																
Audio Mixer Mode	The configuration file <code>/kernel/drv/audiocs.conf</code> is used to configure the <code>audiocs</code> driver so that the audio mixer is enabled or disabled. See the <code>mixer(7I)</code> manual page for details. The audio mixer's mode may be changed at any time using the <code>mixerctl(1)</code> or <code>sdtaudiocontrol(1)</code> applications.																																																
Audio Data Formats	The <code>audiocs</code> device supports 8-bit u-law and A-law, 8-bit and 16-bit linear encodings in mono and stereo. With the mixer enabled, a continuous range of sample rates from 5510 to 48000 Hz is supported. With the mixer disabled, the following sample rates are supported: 5510, 6620, 8000, 9600, 11025, 16000, 18900, 22050, 27420, 32000, 33075, 37800, 4410, and 48000 Hz. When the audio mixer is disabled and the device is open for simultaneous play and record, the input and output data formats must match.																																																
Sample Granularity	At any given time the reported input and output sample counts will vary from the actual sample count by no more than the size of the buffers the <code>audiocs</code> driver is transferring. In general, programs should not rely on the absolute accuracy of the <code>play.samples</code> and <code>record.samples</code> fields of the <code>audio_info</code> structure.																																																

Interrupt Rate	The driver determines how often play and record interrupts should occur. For playing audio, this determines how often and how much audio is requested from the audio mixer. The impact of interrupts on recording is minimal. However, if a very small read buffer size is set, the record interrupt rate should be increased to prevent the buffer from overflowing. The play and record interrupt rates are tunable in the <code>/kernel/drv/audiocs.conf</code> file.								
Audio Status Change Notification	As described in the <code>audio(7I)</code> and <code>mixer(7I)</code> man pages, it is possible to request asynchronous notification of changes in the state of an audio device.								
ERRORS	audiocs errors are described in the <code>audio(7I)</code> man pages.								
FILES	<p><code>/dev/audio</code> Symbolic link to the system's primary audio device. (not necessarily an audiocs audio device).</p> <p><code>/dev/audioctrl</code> Control device for the above audio device.</p> <p><code>/dev/sound/0</code> Represents the first audio device on the system. (not necessarily an audiocs audio device).</p> <p><code>/dev/sound/0ctrl</code> Audio control for above device.</p> <p><code>/usr/share/audio/samples</code> Audio sample files.</p> <p><code>/kernel/drv/sparcv9/audiocs</code> 64-bit audiocs driver.</p> <p><code>/kernel/drv/audiocs.conf</code> Audiocs driver configuration file.</p>								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>SPARC</td> </tr> <tr> <td>Availability</td> <td>SUNWaudd, SUNWauda</td> </tr> <tr> <td>Stability level</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC	Availability	SUNWaudd, SUNWauda	Stability level	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Architecture	SPARC								
Availability	SUNWaudd, SUNWauda								
Stability level	Evolving								
SEE ALSO	<p><code>mixerctl(1)</code>, <code>sdaudiocontrol(1)</code>, <code>ioctl(2)</code>, <code>attributes(5)</code>, <code>audio(7I)</code>, <code>mixer(7I)</code>, <code>streamio(7I)</code></p> <p>Crystal Semiconductor, Inc. CS4231 Data Sheet</p> <p>http://www.sun.com/io</p>								
DIAGNOSTICS	In addition to being logged, the following messages may appear on the system console:								

audiocs(7D)

attach () play interrupt rate set too low.

The play interrupt rate specified in `audiocs.conf` is set too low. It is being reset to the rate specified in the message. Update `audiocs.conf` to a higher play interrupt rate.

attach() play interrupt rate set too high.

The play interrupt rate specified in `audiocs.conf` is set too high. It is being reset to the rate specified in the message. Update `audiocs.conf` to a lower play interrupt rate.

attach() record interrupt rate set too low.

The record interrupt rate specified in `audiocs.conf` is set too low. It is being reset to the rate specified in the message. Update `audiocs.conf` to a higher record interrupt rate.

attach() record interrupt rate set too high

The record interrupt rate specified in `audiocs.conf` is set too high. It is being reset to the rate specified in the message. Update `audiocs.conf` to a lower record interrupt rate.

NAME	audioens – Ensoniq 1371/1373 and Creative Labs 5880 driver support
DESCRIPTION	The audioens driver provides support for the Ensoniq 1371/1373 and Creative Labs 5880 audio controllers. Ensoniq 1371/1373 and Creative Labs 5880 chips are found on add-in PCI cards commonly identified as Audio PCI, CT4740, CT4810, CT5803, CT5808, and ES1371, and on some IA motherboards.
PROGRAM INTERFACE	This interface is described in the mixer(7I) and audio(7I) man pages.
Driver Versions	<p>Applications that open /dev/audio may use the AUDIO_GETDEV ioctl(2) to determine which audio device is being used. The audioens driver returns the string SUNW, audioens in the name field of the audio_device structure. The version field contains a letter (defined in the table below) and the config field contains the string onboard1.</p> <p>The driver works on Sun architectures that support PCI slots, including Sparc and IA.</p> <p>Different cards provide different input and output connectors and may internally hardwire these connectors. SpeakerOut implies an internal audio amplifier and LineOut implies the output is to go to an amplified speaker. The audioens driver routes output data to both SpeakerOut and LineOut. The audioens driver supports input selection from LineIn or Microphone.</p> <p>Platform Type: Sun4u/IA</p> <p>Version: a</p> <p>Line Out: Yes</p> <p>Headphone: No</p> <p>Int. Speaker: Yes</p> <p>Line In: Yes</p> <p>Microphone: Yes</p> <p>CD-ROM: No</p>
Audio Mixer Mode	The configuration file /kernel/drv/audioens.conf is used to configure the audioens driver and determines whether the audio mixer is enabled or disabled. See the mixer(7I) manual page for details. You can change the audio mixer mode at any time using the sdtaudiocontrol(1) or mixerctl(1) applications.
Audio Data Formats	<p>The audioens audio device supports the audio data formats shown below. When the audio mixer is disabled and the device is opened for simultaneous play and record, the input and output data formats may differ. Some sample rates are supported when the mixer is disabled (D) that are not supported when the mixer is enabled (E), due to the overly high computational overhead for sample rate conversion.</p> <p>Supported data formats.</p>

audioens(7D)

Sample Rate	Encoding	Precision	Channels	Mode
5510 Hz	u-Law or A-Law	8	1 or 2	D only
6620 Hz	u-Law or A-Law	8	1 or 2	D only
8000 Hz	u-law or A-law	8	1 or 2	E and D
9600 Hz	u-law or A-law	8	1 or 2	E and D
11025 Hz	u-law or A-law	8	1 or 2	E and D
16000 Hz	u-law or A-law	8	1 or 2	E and D
18900 Hz	u-law or A-law	8	1 or 2	E and D
22050 Hz	u-law or A-law	8	1 or 2	E and D
27420 Hz	u-law or A-law	8	1 or 2	D only
32000 Hz	u-law or A-law	8	1 or 2	E and D
33075 Hz	u-law or A-law	8	1 or 2	E and D
37800 Hz	u-law or A-law	8	1 or 2	E and D
44100 Hz	u-law or A-law	8	1 or 2	E and D
48000 Hz	u-law or A-law	8	1 or 2	E and D
5510 Hz	linear	16	1 or 2	D only
6620 Hz	linear	16	1 or 2	D only
8000 Hz	linear	16	1 or 2	E and D
9600 Hz	linear	16	1 or 2	E and D
11025 Hz	linear	16	1 or 2	E and D
16000 Hz	linear	16	1 or 2	E and D
18900 Hz	linear	16	1 or 2	E and D
22050 Hz	linear	16	1 or 2	E and D
27420 Hz	linear	16	1 or 2	D only
32000 Hz	linear	16	1 or 2	E and D
33075 Hz	linear	16	1 or 2	E and D
37800 Hz	linear	16	1 or 2	E and D
44100 Hz	linear	16	1 or 2	E and D
48000 Hz	linear	16	1 or 2	E and D

Sample Granularity	Because the audioens device manipulates buffers of audio data, the reported input and output sample counts will vary at any given time from the actual sample count by no more than the size of the buffers the audioens driver is transferring. In general, programs should not rely on the absolute accuracy of the <code>play.samples</code> and <code>record.samples</code> fields of the <code>audio_info</code> structure. (See audio(7I)).						
Interrupt Rate	The driver determines how often play and record interrupts take place. For playing audio, the driver determines how often and how much audio is requested from the audio mixer. The impact of interrupts on recording is minimal. However, if a very small read buffer size is set, the record interrupt rate should be increased to prevent the buffer from overflowing. You tune the play and record interrupt rates using the <code>/kernel/drv/audioens.conf</code> file.						
Audio Status Change Notification	As described in the mixer(7I) and audio(7I) man pages, you can request asynchronous notification of changes in the state of an audio device.						
ERRORS	Errors are defined in the audio(7I) man page.						
FILES	<p><code>/dev/audio</code> Symbolic link to the system's primary audio device. (Not necessarily an audioens audio device).</p> <p><code>/dev/audioctrl</code> Control device for the primary audio device.</p> <p><code>/dev/sound/0</code> Represents the first audio device on the system. (Not necessarily an audioens audio device).</p> <p><code>/dev/sound/0ctl</code> Audio control for <code>/dev/sound/0</code>.</p> <p><code>/usr/share/audio/samples</code> Audio sample files.</p> <p><code>/kernel/drv/audioens</code> 32-bit audioens driver</p> <p><code>/kernel/drv/audioens.conf</code> Audioens driver configuration file</p>						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>PCI-based systems</td> </tr> <tr> <td>Availability</td> <td>SUNWaudd, SUNWauda</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	PCI-based systems	Availability	SUNWaudd, SUNWauda
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	PCI-based systems						
Availability	SUNWaudd, SUNWauda						
SEE ALSO	mixerctl(1) , sdtaudiocontrol(1) , ioctl(2) , attributes(5) , audio(7I) , mixer(7I) , streamio(7I)						
	<i>Creative Technology Ltd ES1371 Specification</i>						

audioens(7D)

Creative Technology Ltd ES1373 Specification

Creative Technology Ltd CT5880 Specification

<http://www.sun.com/io>

DIAGNOSTICS

In addition to being logged, the following messages may appear on the system console.

`init_state() play interrupt rate set too low`

The play interrupt rate specified in `audioens.conf` is set too low. It has been reset to the rate specified in the message. Update `audioens.conf` to a higher play interrupt rate.

`init_state() play interrupt rate set too high`

The play interrupt rate set in `audioens.conf` is set too high. It has been reset to the rate specified in the message. Update `audioens.conf` to a lower play interrupt rate.

`init_state() record interrupt rate set too low`

The record interrupt rate specified in `audioens.conf` is set too low. It has been reset to the rate specified in the message. Update `audioens.conf` to a higher record interrupt rate.

`init_state() record interrupt rate set too high`

The record interrupt rate specified in `audioens.conf` is set too high. It is being reset to the rate specified in the message. Update `audioens.conf` to a lower record interrupt rate.

NAME	audio_support – audio driver support routines and interface				
SYNOPSIS	<pre>#include <sys/audio.h></pre>				
DESCRIPTION	The audio support module supports audio drivers that use the new audio driver architecture. It also provides a limited number of <code>ioctl(2)</code> functions for application programmers.				
DATA STRUCTURES	The following data structures are defined to manage the different audio device types and channels.				
Device Types	<p>The following enumeration lists a number of generic device types.</p> <pre>typedef enum { UNDEFINED, AUDIO, AUDIOCTL, USER1, USER2, USER3 } audio_device_type_e;</pre> <p>Currently, the Solaris audio mixer implements only the AUDIO and AUDIOCTL audio device types. See the mixer(7I) manual page for details. The USER1, USER2 and USER3 device types allow third parties to write audio personality modules of their own.</p>				
Channel Structure	<p>This structure is used to get and set state information on individual channels.</p> <pre>struct audio_channel { pid_t pid; /* application's process ID */ uint_t ch_number; /* device channel */ audio_device_type_e dev_type; /* device type */ uint_t info_size; /* size of channel's */ /* info structure */ void *info; /* channel state information */ }; typedef struct audio_channel audio_channel_t;</pre> <p>The <code>ch_number</code> must specify the specific channel number to get or set. When the <code>ioctl()</code> returns, the <code>pid</code> contains the process ID of the process that has that channel open and <code>dev_type</code> will contain the device type. If <code>pid</code> is 0 (zero), then the channel is not open. The pointer <code>*info</code> must point to a buffer large enough to hold whatever audio device-related state structure that may be returned. Currently, only the <code>audio_info_t</code> structure is returned. See the audio(7I) and mixer(7I) man pages for more information.</p>				
IOCTLS	<p>The audio support driver provides the following <code>ioctl()</code>:</p> <table border="0"> <tr> <td style="vertical-align: top;">AUDIO_GET_CH_NUMBER</td> <td>This <code>ioctl()</code> returns the channel number pointed to by the file descriptor. It is returned in the integer pointer of the <code>ioctl()</code> argument.</td> </tr> <tr> <td style="vertical-align: top;">AUDIO_GET_CH_TYPE</td> <td>This <code>ioctl()</code> returns the type of channel the process has open. It is returned in the <code>audio_device_type_e</code> enumeration pointer</td> </tr> </table>	AUDIO_GET_CH_NUMBER	This <code>ioctl()</code> returns the channel number pointed to by the file descriptor. It is returned in the integer pointer of the <code>ioctl()</code> argument.	AUDIO_GET_CH_TYPE	This <code>ioctl()</code> returns the type of channel the process has open. It is returned in the <code>audio_device_type_e</code> enumeration pointer
AUDIO_GET_CH_NUMBER	This <code>ioctl()</code> returns the channel number pointed to by the file descriptor. It is returned in the integer pointer of the <code>ioctl()</code> argument.				
AUDIO_GET_CH_TYPE	This <code>ioctl()</code> returns the type of channel the process has open. It is returned in the <code>audio_device_type_e</code> enumeration pointer				

audio_support(7I)

of the `ioctl()` argument.

`AUDIO_GET_NUM_CHS`

This `ioctl()` returns the number of channels the device supports. It is returned in the integer pointer of the `ioctl()` argument.

MACROS

The following macro can be used to initialize data structures. The established convention is that the state corresponding to a field set to -1 will not be modified.

```
AUDIO_INIT(I, S)
```

Where `I` is a pointer to an `info` structure and `S` is the size of that structure.

The following code segment demonstrates how to use this macro:

```
audio_info_t    info;

AUDIO_INIT(&info, sizeof(info));
info.play.port = AUDIO_HEADPHONE;
err = ioctl(audio_fd, AUDIO_SETINFO, &info);
```

ERRORS

EINVAL The `ioctl()` is invalid for this file descriptor. The `audio_channel_t` structure's `info` pointer does not point to a buffer, or the `ch_number` is bad.

ENOMEM The `ioctl()` failed due to lack of memory.

FILES

`/usr/share/audio/samples`
Audio sample files

ATTRIBUTES

See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86
Availability	SUNWaud, SUNWaudx, SUNWaudh
Stability Level	Evolving

SEE ALSO

`ioctl(2)`, `attributes(5)`, [audio\(7I\)](#), [mixer\(7I\)](#), [streamio\(7I\)](#)

FUTURE DIRECTIONS

Over time, additional audio personality modules will be added. Audio application programmers are encouraged to review this man page with each Solaris release for new audio personality modules.

NAME	audiots – Acer Laboratories Inc. M5451 audio processor interface																
DESCRIPTION	The <code>audiots</code> device uses the ALI M5451 audio processor and an AC-97 Codec to implement the audio device interface.																
APPLICATION PROGRAM INTERFACES	This interface is described in the <code>mixer(7I)</code> and <code>audio(7I)</code> man pages.																
Driver Version	Applications that open <code>/dev/audio</code> may use the <code>AUDIO_GETDEV ioctl(2)</code> to determine which audio device is being used. The <code>audiots</code> driver will return the string <code>SUNW, audiots</code> in the <code>name</code> field of the <code>audio_device</code> structure. The <code>version</code> field will contain a letter (defined in the table below) and the <code>config</code> field will contain the string <code>onboard1</code> .																
	<table border="1"> <thead> <tr> <th>Platform</th> <th>Ver.</th> <th>Out</th> <th>H. Phone</th> <th>Spkr.</th> <th>In</th> <th>Mic.</th> <th>CD</th> </tr> </thead> <tbody> <tr> <td>Sun Blade 100</td> <td>a</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>Y</td> <td>N</td> </tr> </tbody> </table>	Platform	Ver.	Out	H. Phone	Spkr.	In	Mic.	CD	Sun Blade 100	a	Y	Y	Y	Y	Y	N
Platform	Ver.	Out	H. Phone	Spkr.	In	Mic.	CD										
Sun Blade 100	a	Y	Y	Y	Y	Y	N										
	Key to the above table: Ver. = Version. Out = Line Out. H. Phone = Head Phone. Spkr = Internal Speaker. In = Line In. Mic. = Microphone. CD = CD-ROM																
	The <code>audiots</code> device provides support for the internal speaker, headphone, line out, line in, and microphone. The <code>play.mod_ports</code> and <code>record.mod_ports</code> fields of the <code>audio_info</code> structure (see <code>audio(7I)</code>) indicate which ports may be manipulated.																
Audio Mixer Mode	The configuration file <code>/kernel/drv/audiots.conf</code> is used to configure the <code>audiots</code> driver so that the audio mixer is enabled or disabled. See the <code>mixer(7I)</code> manual page for details. The audio mixer mode may be changed at any time using the <code>sdaudiocontrol(1)</code> or <code>mixerctl(1)</code> commands.																
Audio Data Formats	The <code>audiots</code> device supports 8-bit μ -law and A-law, 8-bit linear and 16-bit linear encodings in mono and stereo. With the mixer enabled, a continuous range of sample rates from 5510 to 48000 Hz is supported. With the mixer disabled, the following sample rates are supported: 5510, 6620, 8000, 9600, 11025, 16000, 18900, 22050, 27420, 32000, 33075, 37800, 4410, and 48000 Hz. When the audio mixer is disabled and the device is opened for simultaneous play and record, the input and out data formats may be different.																
Sample Granularity	Because the <code>audiots</code> device manipulates buffers of audio data, the reported input and output sample counts will vary at any given time from the actual sample count by no more than the size of the buffers the <code>audiots</code> driver is transferring. In general, programs should not rely on the absolute accuracy of the <code>play.samples</code> and <code>record.samples</code> fields of the <code>audio_info</code> structure.																

audiots(7D)

Interrupt Rate The driver determines how often play and record interrupts should take place. For playing audio, this determines how often and how much audio is requested from the audio mixer. The impact of interrupts on recording is minimal. However, if a very small read buffer size is set, the record interrupt rate should be increased to prevent the buffer from overflowing. The play and record interrupt rates are tunable in the `/kernel/drv/audiots.conf` file.

Audio Status Change Notification As described in the `audio(7I)` and `mixer(7I)` man pages, it is possible to request asynchronous notification of changes in the state of an audio device.

ERRORS `audiots` errors are described in the `audio(7I)` man page.

FILES

<code>/dev/audio</code>	Symbolic link to the system's primary audio device. (Not necessarily an <code>audiots</code> audio device).
<code>/dev/audioctrl</code>	Control device for the primary audio device.
<code>/dev/sound/0</code>	Represents the first audio device on the system. (Not necessarily an <code>audiots</code> audio device).
<code>/dev/sound/0ctrl</code>	Audio control for <code>/dev/sound/0</code> .
<code>/usr/share/audio/samples</code>	Audio sample files.
<code>/kernel/drv/sparcv9/audiots</code>	64-bit <code>audiots</code> driver
<code>/kernel/drv/audiots.conf</code>	<code>audiots</code> driver configuration file.

ATTRIBUTES See `attributes(5)` for a descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Availability	SUNWaudd, SUNWauddx, SUNWauda
Stability level	Evolving

SEE ALSO `mixerctl(1)`, `sdtaudiocontrol(1)`, `ioctl(2)`, `attributes(5)`, `audio(7I)`, `mixer(7I)`, `streamio(7I)`

Acer Laboratories Inc. M5451 PCI Audio Processor Technical Specification

<http://www.sun.com/io>

DIAGNOSTICS In addition to being logged, the following messages may appear on the system console:

`init_state()` play interrupt rate set too low

The play interrupt rate in `audiots.conf` is set too low. It has been reset to the rate specified in the message. Update `audiots.conf` to a higher play interrupt rate.

`init_state()` play interrupt rate set too high

The play interrupt rate set in `audiots.conf` is set too high. It has been reset to the rate specified in the message. Update `audiots.conf` to a lower play interrupt rate.

`init_state()` record interrupt rate set too low

The record interrupt rate in `audiots.conf` is set too low. It has been reset to the rate specified in the message. Update `audiots.conf` to a higher record interrupt rate.

`init_state()` record interrupt rate set too high

The record interrupt rate in `audiots.conf` is set too high. It has been reset to the rate specified in the message. Update `audiots.conf` to a lower record interrupt rate.

av1394(7D)

NAME av1394 – 1394 audio/video driver

SYNOPSIS unit@GUID

DESCRIPTION The av1394 driver is an *IEEE 1394* compliant target driver that supports the *IEC 61883* Consumer Audio/Video Equipment - Digital Interface standard. The driver is used to receive and transmit isochronous data streams in the common isochronous packet (CIP) format, as well as asynchronous function control protocol (FCP) frames. The driver also supports connection management procedures (CMP).

DEVICE SPECIAL FILES

/dev/av/N/async
Device node for asynchronous data

/dev/av/N/isochn
Device node for isochronous data

/kernel/drv/amd64/av1394
64-bit kernel module for 64-bit x86 platform

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86, PCI-based systems
Availability	SUNWav1394

SEE ALSO [hci1394\(7D\)](#)

IEEE Std 1394-1995 Standard for a High Performance Serial Bus

IEC 61883 Consumer Audio/Video Equipment - Digital Interface

NAME	bbc_beep – Platform-dependent Beep driver for BBC-based hardware.						
SYNOPSIS	beep@unit-address						
DESCRIPTION	The <code>bbc_beep</code> driver generates beeps on platforms (including Sun Blade 1000) that use BBC-based registers and USB keyboards. When the <code>KIOCCMD</code> ioctl is issued to the USB keyboard module (see <code>usbkbm(7M)</code>) with command <code>KBD_CMD_BELL/KBD_CMD_NOBELL</code> , <code>usbkbm(7M)</code> passes the request to the <code>bbc_beep</code> driver to turn the beep on and off, respectively.						
FILES	/platform/sun4u/kernel/drv/sparcv9/bbc_beep 64-bit ELF kernel driver						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>BBC-based SPARC</td> </tr> <tr> <td>Availability</td> <td>SUNWcarx.u</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	BBC-based SPARC	Availability	SUNWcarx.u
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	BBC-based SPARC						
Availability	SUNWcarx.u						
SEE ALSO	<code>kbd(1)</code> , <code>attributes(5)</code> , <code>grbeep(7D)</code> , <code>kb(7M)</code> , <code>usbkbm(7M)</code> <i>Writing Device Drivers</i>						
DIAGNOSTICS	None						

bd(7M)

NAME	bd – SunButtons and SunDials STREAMS module
SYNOPSIS	<code>open("/dev/bd", O_RDWR)</code>
DESCRIPTION	<p>The bd STREAMS module processes the byte streams generated by the SunButtons buttonbox and SunDials dialbox. The buttonbox generates a stream of bytes that encode the identity and state transition of the buttons. The dialbox generates a stream of bytes that encode the identity of the dials and the amount by which they are turned. Both of these streams are merged together when a host has both a buttonbox and a dialbox in use at the same time.</p> <p>SunButtons reports the button number and up/down status encoded into a one byte message. Byte values from 0xc0 to 0xdf indicate a transition to button down. To obtain the button number, subtract 0xc0 from the byte value. Byte values from 0xe0 to 0xff indicate a transition to button up. To obtain the button number, subtract 0xe0 from the byte value.</p> <p>Each dial sample in the byte stream consists of three bytes. The first byte identifies which dial was turned and the next two bytes return the delta in signed binary format. When bound to an application using the window system, Virtual User Input Device (“VUID”) events are generated. An event from a dial is constrained to lie between 0x80 and 0x87.</p> <p>A stream with the bd pushed streams module configured in it can emit <code>firm_events</code> as specified by the protocol of a VUID. bd understands the <code>VUIDSFORMAT</code> and <code>VUIDGFORMAT</code> ioctls (see reference below), as defined in <code>/usr/include/sys/bdio.h</code> and <code>\$OPENWINHOME/include/xview/win_event.h</code>. All other <code>ioctl()</code> requests are passed downstream.</p> <p>The bd streams module sets the parameters of the serial port when it is first opened. No <code>termio(7I)</code> <code>ioctl()</code> requests should be performed on a bd STREAMS module, as bd expects the device parameters to remain as it set them.</p>
IOCTLS	<p><code>VUIDSFORMAT</code> <code>VUIDGFORMAT</code> These are standard VUID ioctls.</p> <p><code>BDIOBUTLITE</code> The bd streams module implements this ioctl to enable processes to manipulate the lights on the buttonbox. The <code>BDIOBUTLITE</code> ioctl must be carried by an <code>I_STR</code> ioctl to the bd module. For an explanation of <code>I_STR</code> see <code>streamio(7I)</code>. The data for the <code>BDIOBUTLITE</code> ioctl is an unsigned integer in which each bit represents the lamp on one button. The macro <code>LED_MAP</code> in <code><sys/bdio.h></code> maps button numbers to appropriate bits. Source code for the demo program <code>x_buttonstest</code> is provided with the buttons and dials package, and may be found in the directory <code>/usr/demo/BUTTONBOX</code>. Look at <code>x_buttonstest.c</code> for an example of how to manipulate the lights on the buttonbox.</p>
FILES	<code>/usr/include/sys/bdio.h</code>

/usr/include/sys/stropts.h

\$OPENWINHOME/share/include/xview/win_event.h

SEE ALSO bdconfig(1M), ioctl(2), x_buttontest(6), x_dialtest(6), streamio(7I),
termio(7I)

SunButtons Installation and Programmers Guide

SunDials Installation and Programmers Guide

WARNINGS The SunDials dial box must be used with a serial port.

bge(7D)

NAME	bge – SUNW,bge Gigabit Ethernet driver for Broadcom BCM57xx
SYNOPSIS	/dev/bge*
DESCRIPTION	<p>The bge Gigabit Ethernet driver is a multi-threaded, loadable, clonable, GLD-based STREAMS driver supporting the Data Link Provider Interface, <code>dlpi(7P)</code>, on Broadcom BCM57xx (BCM5703/5704, and BCM5700/5701/5705/5782 on x86) Gigabit Ethernet controllers fitted to the system motherboard. With the exception of BCM5700/BCM5701/BCM5704S, these devices incorporate both MAC and PHY functions and provide three-speed (copper) Ethernet operation on the RJ-45 connectors. (BCM5700/BCM5701/BCM5704S do not have a PHY integrated into the MAC chipset.)</p> <p>The bge driver functions include controller initialization, frame transmit and receive, promiscuous and multicast support, and error recovery and reporting.</p> <p>The bge driver and hardware support auto-negotiation, a protocol specified by the 1000 Base-T standard. Auto-negotiation allows each device to advertise its capabilities and discover those of its peer (link partner). The highest common denominator supported by both link partners is automatically selected, yielding the greatest available throughput, while requiring no manual configuration. The bge driver also allows you to configure the advertised capabilities to less than the maximum (where the full speed of the interface is not required), or to force a specific mode of operation, irrespective of the link partner's advertised capabilities.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device, <code>/dev/bge</code>, is used to access all BCM57xx devices (BCM5703/5704, and BCM5700/5701/5705/5782 on x86) fitted to the system motherboard.</p> <p>The bge driver is dependent on <code>/kernel/misc/gld</code>, a loadable kernel module that provides the bge driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on supported primitives.</p> <p>You must send an explicit <code>DL_ATTACH_REQ</code> message to associate the opened stream with a particular device (PPA). The PPA ID is interpreted as an unsigned integer data type and indicates the corresponding device instance (unit) number. The driver returns an error (<code>DL_ERROR_ACK</code>) if the PPA field value does not correspond to a valid device instance number for the system. The device is initialized on first attach and de-initialized (stopped) at last detach.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to a <code>DL_INFO_REQ</code> are:</p> <ul style="list-style-type: none">■ Maximum SDU is 1500 (<code>ETHERMTU</code> - defined in <code><sys/ethernet.h></code>).■ Minimum SDU is 0.■ DLSAP address length is 8.■ MAC type is <code>DL_ETHER</code>.

- SAP length value is -2, meaning the physical address component is followed immediately by a 2-byte SAP component within the DLSAP address.
- Broadcast address value is the Ethernet/IEEE broadcast address (FF:FF:FF:FF:FF:FF).

Once in the DL_ATTACHED state, you must send a DL_BIND_REQ to associate a particular Service Access Point (SAP) with the stream.

CONFIGURATION

By default, the bge driver performs auto-negotiation to select the link speed and mode. Link speed and mode can be any one of the following, (as described in the IEEE803.2 standard):

- 1000 Mbps, full-duplex
- 1000 Mbps, half-duplex
- 100 Mbps, full-duplex
- 100 Mbps, half-duplex
- 10 Mbps, full-duplex
- 10 Mbps, half-duplex

The auto-negotiation protocol automatically selects:

- Speed (1000 Mbps, 100 Mbps, or 10 Mbps)
- Operation mode (full-duplex or half-duplex)

as the highest common denominator supported by both link partners. Because the bge device supports all modes, the effect is to select the highest throughput mode supported by the other device.

Alternatively, you can set the capabilities advertised by the bge device using `ndd(1M)`. The driver supports a number of parameters whose names begin with `adv_` (see below). Each of these parameters contains a boolean value that determines whether the device advertises that mode of operation. The `adv_pause_cap` indicates whether half/full duplex pause is advertised to link partner. And the `adv_asym_pause_cap` can be set to advertise to link partner that asymmetric pause is desired. In addition, bge uses `adv_100T4_cap` to advertise its 100T4 capability. The `adv_autoneg_cap` parameter controls whether autonegotiation is performed. If `adv_autoneg_cap` is set to 0, the driver forces the mode of operation selected by the first non-zero parameter in priority order as listed below:

```

                                (highest priority/greatest throughput)
adv_1000fdx_cap      1000Mbps full duplex
adv_1000hdx_cap     1000Mbps half duplex
adv_100fdx_cap      100Mbps full duplex
adv_100hdx_cap      100Mbps half duplex
adv_10fdx_cap       10Mbps full duplex
adv_10hdx_cap       10Mbps half duplex
                                (lowest priority/least throughput)
```

For example, to prevent the device 'bge2' from advertising gigabit capabilities, enter (as super-user):

bge(7D)

```
# ndd -set /dev/bge2 adv_1000hdx_cap 0
# ndd -set /dev/bge2 adv_1000fdx_cap 0
```

All capabilities default to enabled. Note that changing any capability parameter causes the link to go down while the link partners renegotiate the link speed/duplex using the newly changed capabilities.

The current settings of the parameters may be found using `ndd -get`. In addition, the driver exports the current state, speed, duplex setting, and working mode of the link via `ndd` parameters (these are read only and may not be changed). For example, to check link state of device `bge0`:

```
# ndd -get /dev/bge0 link_status
1
# ndd -get /dev/bge0 link_speed
100
# ndd -get /dev/bge0 link_duplex
2
# ndd -get /dev/bge0 link_rx_pause
1
# ndd -get /dev/bge0 link_tx_pause
1
```

The output above indicates that the link is up and running at 100Mbps full-duplex with its rx/tx direction pause capability. In addition, the driver exports its working mode by `loop_mode`. If it is set to 0, the loopback mode is disabled.

FILES

`/dev/bge*`
Character special device.

`/kernel/drv/sparcv9/bge`
SPARC bge driver binary.

`/kernel/drv/amd64/bge`
kernel module for 64-bit x86 platform

`/platform/i86pc/kernel/drv/bge`
32-bit x86 bge driver binary.

`/platform/i86pc/kernel/drv/amd64/bge`
64-bit x86 bge driver binary.

ATTRIBUTES

See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86

SEE ALSO

`attributes(5)`, `gld(7D)`, `streamio(7I)`, `dlpi(7P)`

Writing Device Drivers

STREAMS Programming Guide

bpp(7D)

NAME	bpp – bi-directional parallel port driver
SYNOPSIS	<code>SUNW, bpp@slot, offset : bppn</code>
DESCRIPTION	The <code>bpp</code> driver provides a general-purpose bi-directional interface to parallel devices. It supports a variety of output (printer) and input (scanner) devices, using programmable timing relationships between the various handshake signals.
APPLICATION PROGRAMMING INTERFACE	The <code>bpp</code> driver is an <i>exclusive-use</i> device. If the device has already been opened, subsequent opens fail with <code>EBUSY</code> .
Default Operation	Each time the <code>bpp</code> device is opened, the default configuration is <code>BPP_ACK_BUSY_HS</code> for read handshake, <code>BPP_ACK_HS</code> for write handshake, 1 microsecond for all setup times and strobe widths, and 60 seconds for both timeouts. This configuration (in the write mode) drives many common personal computer parallel printers with Centronics-type interfaces. The application should use the <code>BPPIOC_SETPARMS</code> ioctl request to configure the <code>bpp</code> for the particular device which is attached, if necessary.
Write Operation	If a failure or error condition occurs during a <code>write(2)</code> , the number of bytes successfully written is returned (short write). Note that <code>errno</code> will not be set. The contents of certain status bits will be captured at the time of the error, and can be retrieved by the application program, using the <code>BPPIOC_GETERR</code> ioctl request. Subsequent <code>write(2)</code> calls may fail with the system error <code>ENXIO</code> if the error condition is not rectified. The captured status information will be overwritten each time an attempted transfer or a <code>BPPIOC_TESTIO</code> ioctl request occurs.
Read Operations	<p>If a failure or error condition occurs during a <code>read(2)</code>, the number of bytes successfully read is returned (short read). Note that <code>errno</code> will not be set. The contents of certain status bits will be captured at the time of the error, and can be retrieved by the application, using the <code>BPPIOC_GETERR</code> ioctl request. Subsequent <code>read(2)</code> calls may fail with <code>ENXIO</code> if the error condition is not rectified. The captured register information will be overwritten each time an attempted transfer or a <code>BPPIOC_TESTIO</code> ioctl request.</p> <p>If the <code>read_handshake</code> element of the <code>bpp_transfer_parms</code> structure (see below) is set to <code>BPP_CLEAR_MEM</code> or <code>BPP_SET_MEM</code>, zeroes or ones, respectively, are written into the user buffer.</p>
Read/Write Operation	When the driver is opened for reading and writing, it is assumed that scanning will take place, as scanners are the only devices supported by this mode. Most scanners require that the <code>SLCT_IN</code> or <code>AFX</code> pin be set to tell the scanner the direction of the transfer. The <code>AFX</code> line is set when the <code>read_handshake</code> element of the <code>bpp_transfer_parms</code> structure is set to <code>BPP_HSCAN_HS</code> , otherwise the <code>SLCT_IN</code> pin is set. Normally, scanning starts by writing a command to the scanner, at which time the pin is set. When the scan data is read back, the pin is reset.
IOCTLS	The following ioctl requests are supported: <code>BPPIOC_SETPARMS</code> Set transfer parameters.

BPPIOC_GETPARMS	<p>The argument is a pointer to a <code>bpp_transfer_parms</code> structure. See below for a description of the elements of this structure. If a parameter is out of range, <code>EINVAL</code> is returned.</p> <p>Get current transfer parameters.</p>
BPPIOC_SETOUTPINS	<p>The argument is a pointer to a <code>bpp_transfer_parms</code> structure. See below for a description of the elements of this structure. If no parameters have been configured since the device was opened, the contents of the structure will be the default conditions of the parameters (see <code>Default Operation</code> above).</p> <p>Set output pin values.</p>
BPPIOC_GETOUTPINS	<p>The argument is a pointer to a <code>bpp_pins</code> structure. See below for a description of the elements of this structure. If a parameter is out of range, <code>EINVAL</code> is returned.</p> <p>Read output pin values. The argument is a pointer to a <code>bpp_pins</code> structure. See below for a description of the elements of this structure.</p>
BPPIOC_GETERR	<p>Get last error status.</p> <p>The argument is a pointer to a <code>bpp_error_status</code> structure. See below for a description of the elements of this structure. This structure indicates the status of all the appropriate status bits at the time of the most recent error condition during a <code>read(2)</code> or <code>write(2)</code> call, or the status of the bits at the most recent <code>BPPIOC_TESTIO</code> ioctl request. Note: The bits in the <code>pin_status</code> element indicate whether the associated pin is active, not the actual polarity. The application can check transfer readiness without attempting another transfer using the <code>BPPIOC_TESTIO</code> ioctl. Note: The <code>timeout_occurred</code> and <code>bus_error</code> fields will never be set by the <code>BPPIOC_TESTIO</code> ioctl, only by an actual failed transfer.</p>
BPPIOC_TESTIO	<p>Test transfer readiness.</p> <p>This command checks to see if a read or write transfer would succeed based on pin status, opened mode, and handshake selected. If a handshake would succeed, 0 is returned. If a transfer would fail, -1 is returned, and <code>errno</code> is set to <code>EIO</code>, and the error status information is captured. The captured status can be retrieved using</p>

bpp(7D)

Transfer Parameters Structure

the BPPIOC_GETERR ioctl call. Note that the timeout_occurred and bus_error fields will never be set by this ioctl.

This structure is defined in <sys/bpp_io.h>.

```
struct bpp_transfer_parms {
    enum handshake_t
        read_handshake; /* parallel port read handshake mode */
    int read_setup_time; /* DSS register - in nanoseconds */
    int read_strobe_width; /* DSW register - in nanoseconds */
    int read_timeout; /*
        * wait this many seconds
        * before aborting a transfer
        */

    enum handshake_t
        write_handshake; /* parallel port write handshake mode */
    int write_setup_time; /* DSS register - in nanoseconds */
    int write_strobe_width; /* DSW register - in nanoseconds */
    int write_timeout; /*
        * wait this many seconds
        * before aborting a transfer
        */
};
/* Values for read_handshake and write_handshake fields */
enum handshake_t {
    BPP_NO_HS, /* no handshake pins */
    BPP_ACK_HS, /* handshake controlled by ACK line */
    BPP_BUSY_HS, /* handshake controlled by BSY line */
    BPP_ACK_BUSY_HS, /*
        * handshake controlled by ACK and BSY lines
        * read_handshake only!
        */
    BPP_XSCAN_HS, /* xerox scanner mode,
        * read_handshake only!
        */
    BPP_HSCAN_HS, /*
        * HP scanjet scanner mode
        * read_handshake only!
        */
    BPP_CLEAR_MEM, /* write 0's to memory,
        * read_handshake only!
        */
    BPP_SET_MEM, /* write 1's to memory,
        * read_handshake only!
        */
    /* The following handshakes are RESERVED. Do not use. */
    BPP_VPRINT_HS, /* valid only in read/write mode */
    BPP_VPLOT_HS /* valid only in read/write mode */
};
```

The `read_setup_time` field controls the time between `dstrb` falling edge to `bsy` rising edge if the `read_handshake` field is set to `BPP_NO_HS` or `BPP_ACK_HS`. It controls the time between `dstrb` falling edge to `ack` rising edge if the `read_handshake` field is set to `BPP_ACK_HS` or `BPP_ACK_BUSY_HS`. It controls the time between `ack` falling edge to `dstrb` rising edge if the `read_handshake` field is set to `BPP_XSCAN_HS`.

The `read_strobe_width` field controls the time between `ack` rising edge and `ack` falling edge if the `read_handshake` field is set to `BPP_NO_HS` or `BPP_ACK_BUSY_HS`. It controls the time between `dstrb` rising edge to `dstrb` falling edge if the `read_handshake` field is set to `BPP_XSCAN_HS`.

The values allowed for the `write_handshake` field are duplicates of the definitions for the `read_handshake` field. Note that some of these handshake definitions are only valid in one mode or the other.

The `write_setup_time` field controls the time between data valid to `dstrb` rising edge for all values of the `write_handshake` field.

The `write_strobe_width` field controls the time between `dstrb` rising edge and `dstrb` falling edge if the `write_handshake` field is not set to `BPP_VPRINT_HS` or `BPP_VPLOT_HS`. It controls the minimum time between `dstrb` rising edge to `dstrb` falling edge if the `write_handshake` field is set to `BPP_VPRINT_HS` or `BPP_VPLOT_HS`.

Transfer Pins Structure

This structure is defined in `<sys/bpp_io.h>`.

```
struct bpp_pins {
    uchar_t  output_reg_pins; /* pins in P_OR register */
    uchar_t  input_reg_pins; /* pins in P_IR register */
};

/* Values for output_reg_pins field */
#define BPP_SLCTIN_PIN 0x01 /* Select in pin */
#define BPP_AFX_PIN    0x02 /* Auto feed pin */
#define BPP_INIT_PIN   0x04 /* Initialize pin */
#define BPP_V1_PIN     0x08 /* reserved pin 1 */
#define BPP_V2_PI      0x10 /* reserved pin 2 */
#define BPP_V3_PIN     0x20 /* reserved pin 3 */
#define BPP_ERR_PIN    0x01 /* Error pin */
#define BPP_SLCT_PIN   0x02 /* Select pin */
#define BPP_PE_PIN     0x04 /* Paper empty pin */
```

Error Pins Structure

This structure is defined in the include file `<sys/bpp_io.h>`.

```
struct bpp_error_status {
    char  timeout_occurred; /* 1 if a timeout occurred */
    char  bus_error; /* 1 if an SBus bus error */
    uchar_t pin_status; /*
        * status of pins which could
        * cause an error
        */
};
```

bpp(7D)

	<pre>/* Values for pin_status field */ #define BPP_ERR_ERR 0x01 /* Error pin active */ #define BPP_SLCT_ERR 0x02 /* Select pin active */ #define BPP_PE_ERR 0x04 /* Paper empty pin active */ #define BPP_SLCTIN_ERR 0x10 /* Select in pin active */ #define BPP_BUSY_ERR 0x40 /* Busy pin active */</pre>	
ERRORS	EBADF	The device is opened for write-only access and a read is attempted, or the device is opened for read-only access and a write is attempted.
	EBUSY	The device has been opened and another open is attempted. An attempt has been made to unload the driver while one of the units is open.
	EINVAL	A <code>BPPIOC_SETPARMS</code> <code>ioctl</code> is attempted with an out of range value in the <code>bpp_transfer_parms</code> structure. A <code>BPPIOC_SETOUTPINS</code> <code>ioctl</code> is attempted with an invalid value in the <code>pins</code> structure. An <code>ioctl</code> is attempted with an invalid value in the command argument. An invalid command argument is received during <code>modload(1M)</code> or <code>modunload(1M)</code> .
	EIO	The driver encountered an SBus bus error when attempting an access. A read or write does not complete properly, due to a peripheral error or a transfer timeout. A <code>BPPIOC_TESTIO</code> <code>ioctl</code> call is attempted while a condition exists which would prevent a transfer (such as a peripheral error).
	ENXIO	The driver has received an open request for a unit for which the attach failed. The driver has received a read or write request for a unit number greater than the number of units available. The driver has received a write request for a unit which has an active peripheral error.
FILES	<code>/dev/bppn</code>	bi-directional parallel port devices
SEE ALSO	<code>ioctl(2)</code> , <code>read(2)</code> , <code>write(2)</code> , <code>sbus(4)</code>	

NAME	bscv, bscbus, i2bsc – Blade support chip interface driver
DESCRIPTION	<p>The <i>bscv</i>, <i>bscbus</i> and <i>i2bsc</i> drivers interface with the Blade support chip used on Sun Microsystem's Blade server products. These drivers provide a conduit for passing control, environmental, cpu signature and event information between Solaris and the Blade support chip.</p> <p>These drivers do not export public interfaces. Instead they make information available via <i>picl</i>, <i>prtdiag</i>, <i>prtfu</i> and related tools. In addition, these drivers log Blade support chip environmental event information into system logs.</p>
FILES	<p>/platform/sun4u/kernel/drv/sparcv9/bscbus 64-bit ELF kernel driver</p> <p>/platform/sun4u/kernel/drv/sparcv9/bscv 64-bit ELF kernel driver</p> <p>/platform/sun4u/kernel/drv/sparcv9/i2bsc 64-bit ELF kernel driver</p> <p>/platform/i86pc/kernel/drv/bscbus 32-bit ELF kernel file (x86 only)</p> <p>/platform/i86pc/kernel/drv/bscv 32-bit ELF kernel file (x86 only)</p>
ATTRIBUTES	See <i>attributes(5)</i> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Limited to systems with Blade Support Chip
Availability	SUNWcar.u, SUNWcar.i

bufmod(7M)

NAME	bufmod – STREAMS Buffer Module
SYNOPSIS	<pre>#include <sys/bufmod.h> ioctl(fd, I_PUSH, "bufmod");</pre>
DESCRIPTION	<p>bufmod is a STREAMS module that buffers incoming messages, reducing the number of system calls and the associated overhead required to read and process them. Although bufmod was originally designed to be used in conjunction with STREAMS-based networking device drivers, the version described here is general purpose so that it can be used anywhere STREAMS input buffering is required.</p>
Read-side Behavior	<p>The behavior of bufmod depends on various parameters and flags that can be set and queried as described below under <code>IOCTLs</code>. bufmod collects incoming <code>M_DATA</code> messages into chunks, passing each chunk upstream when the chunk becomes full or the current read timeout expires. It optionally converts <code>M_PROTO</code> messages to <code>M_DATA</code> and adds them to chunks as well. It also optionally adds to each message a header containing a timestamp, and a cumulative count of messages dropped on the stream read side due to resource exhaustion or flow control. The default settings of bufmod allow it to drop messages when flow control sets in or resources are exhausted; disabling headers and explicitly requesting no drops makes bufmod pass all messages through. Finally, bufmod is capable of truncating upstream messages to a fixed, programmable length.</p> <p>When a message arrives, bufmod processes it in several steps. The following paragraphs discuss each step in turn.</p> <p>Upon receiving a message from below, if the <code>SB_NO_HEADER</code> flag is not set, bufmod immediately timestamps it and saves the current time value for later insertion in the header described below.</p> <p>Next, if <code>SB_NO_PROTO_CVT</code> is not set, bufmod converts all leading <code>M_PROTO</code> blocks in the message to <code>M_DATA</code> blocks, altering only the message type field and leaving the contents alone.</p> <p>It then truncates the message to the current <i>snapshot length</i>, which is set with the <code>SBIOCSSNAP</code> ioctl described below.</p> <p>Afterwards, if <code>SB_NO_HEADER</code> is not set, bufmod prepends a header to the converted message. This header is defined as follows.</p> <pre>struct sb_hdr { uint_t sbh_origlen; uint_t sbh_msglen; uint_t sbh_totlen; uint_t sbh_drops; #ifdef _LP64 defined(_I32LPx) struct timeval32 sbh_timestamp; #else struct timeval sbh_timestamp; #endif /* !_LP64 */ };</pre>

The `sbh_origlen` field gives the message's original length before truncation in bytes. The `sbh_msglen` field gives the length in bytes of the message after the truncation has been done. `sbh_totlen` gives the distance in bytes from the start of the truncated message in the current chunk (described below) to the start of the next message in the chunk; the value reflects any padding necessary to insure correct data alignment for the host machine and includes the length of the header itself. `sbh_drops` reports the cumulative number of input messages that this instance of `bufmod` has dropped due to flow control or resource exhaustion. In the current implementation message dropping due to flow control can occur only if the `SB_NO_DROPS` flag is not set. (Note: this accounts only for events occurring within `bufmod`, and does not count messages dropped by downstream or by upstream modules.) The `sbh_timestamp` field contains the message arrival time expressed as a `struct timeval`.

After preparing a message, `bufmod` attempts to add it to the end of the current chunk, using the chunk size and timeout values to govern the addition. The chunk size and timeout values are set and inspected using the `ioctl()` calls described below. If adding the new message would make the current chunk grow larger than the chunk size, `bufmod` closes off the current chunk, passing it up to the next module in line, and starts a new chunk. If adding the message would still make the new chunk overflow, the module passes it upward in an over-size chunk of its own. Otherwise, the module concatenates the message to the end of the current chunk.

To ensure that messages do not languish forever in an accumulating chunk, `bufmod` maintains a read timeout. Whenever this timeout expires, the module closes off the current chunk and passes it upward. The module restarts the timeout period when it receives a read side data message and a timeout is not currently active. These two rules insure that `bufmod` minimizes the number of chunks it produces during periods of intense message activity and that it periodically disposes of all messages during slack intervals, but avoids any timeout overhead when there is no activity.

`bufmod` handles other message types as follows. Upon receiving an `M_FLUSH` message specifying that the read queue be flushed, the module clears the currently accumulating chunk and passes the message on to the module or driver above. (Note: `bufmod` uses zero length `M_CTL` messages for internal synchronization and does not pass them through.) `bufmod` passes all other messages through unaltered to its upper neighbor, maintaining message order for non high priority messages by passing up any accumulated chunk first.

If the `SB_DEFER_CHUNK` flag is set, buffering does not begin until the second message is received within the timeout window.

If the `SB_SEND_ON_WRITE` flag is set, `bufmod` passes up the read side any buffered data when a message is received on the write side. `SB_SEND_ON_WRITE` and `SB_DEFER_CHUNK` are often used together.

Write-side Behavior

`bufmod` intercepts `M_IOCTL` messages for the `ioctls` described below. The module passes all other messages through unaltered to its lower neighbor. If `SB_SEND_ON_WRITE` is set, message arrival on the writer side suffices to close and transmit the current read side chunk.

bufmod(7M)

IOCTLS	bufmod responds to the following <code>ioctl</code> s.	
<code>SBIOCSTIME</code>	Set the read timeout value to the value referred to by the <code>struct timeval</code> pointer given as argument. Setting the timeout value to zero has the side-effect of forcing the chunk size to zero as well, so that the module will pass all incoming messages upward immediately upon arrival. Negative values are rejected with an <code>EINVAL</code> error.	
<code>SBIOCGTIME</code>	Return the read timeout in the <code>struct timeval</code> pointed to by the argument. If the timeout has been cleared with the <code>SBIOCCTIME</code> <code>ioctl</code> , return with an <code>ERANGE</code> error.	
<code>SBIOCCTIME</code>	Clear the read timeout, effectively setting its value to infinity. This results in no timeouts being active and the chunk being delivered when it is full.	
<code>SBIOCSCHUNK</code>	Set the chunk size to the value referred to by the <code>uint_t</code> pointer given as argument. See NOTES for a description of effect on stream head high water mark.	
<code>SBIOCGCHUNK</code>	Return the chunk size in the <code>uint_t</code> pointed to by the argument.	
<code>SBIOCSSNAP</code>	Set the current snapshot length to the value given in the <code>uint_t</code> pointed to by the <code>ioctl</code> 's final argument. <code>bufmod</code> interprets a snapshot length value of zero as meaning infinity, so it will not alter the message. See NOTES for a description of effect on stream head high water mark.	
<code>SBIOCGSNAP</code>	Returns the current snapshot length in the <code>uint_t</code> pointed to by the <code>ioctl</code> 's final argument.	
<code>SBIOCFLAGS</code>	Set the current flags to the value given in the <code>uint_t</code> pointed to by the <code>ioctl</code> 's final argument. Possible values are a combination of the following.	
	<code>SB_SEND_ON_WRITE</code>	Transmit the read side chunk on arrival of a message on the write side.
	<code>SB_NO_HEADER</code>	Do not add headers to read side messages.
	<code>SB_NO_DROPS</code>	Do not drop messages due to flow control upstream.
	<code>SB_NO_PROTO_CVT</code>	Do not convert <code>M_PROTO</code> messages into <code>M_DATA</code> .
	<code>SB_DEFER_CHUNK</code>	Begin buffering on arrival of the second read side message in a timeout interval.

`SBIOCGFLAGS` Returns the current flags in the `uint_t` pointed to by the `ioctl`'s final argument.

SEE ALSO [dlpi\(7P\)](#), [pfmod\(7M\)](#)

NOTES Older versions of `bufmod` did not support the behavioral flexibility controlled by the `SBIOCGFLAGS` `ioctl`. Applications that wish to take advantage of this flexibility can guard themselves against old versions of the module by invoking the `SBIOCGFLAGS` `ioctl` and checking for an `EINVAL` error return.

When buffering is enabled by issuing an `SBIOCSCHUNK` `ioctl` to set the chunk size to a non zero value, `bufmod` sends a `SETOPTS` message to adjust the stream head high and low water marks to accommodate the chunked messages.

When buffering is disabled by setting the chunk size to zero, message truncation can have a significant influence on data traffic at the stream head and therefore the stream head high and low water marks are adjusted to new values appropriate for the smaller truncated message sizes.

BUGS `bufmod` does not defend itself against allocation failures, so that it is possible, although very unlikely, for the stream head to use inappropriate high and low water marks after the chunk size or snapshot length have changed.

cadp160(7D)

NAME	cadp160 – Adaptec Ultra160 SCSI host bus adapter driver				
SYNOPSIS	scsi@unit-address				
DESCRIPTION	<p>The cadp160 host bus adapter driver is a SCSI-compliant nexus driver that supports the following Adaptec Ultra160 SCSI devices:</p> <ul style="list-style-type: none">■ Adapters: 39160, 29160, 29160N, 29160LP■ Chips: AIC-7892B1, AIC-7899A, AIC-7899B2 <p>The cadp160 driver supports standard functions provided by the SCSI interface including tagged and untagged queuing, wide, fast and ultra SCSI, and auto request sense. The cadp160 driver does not support linked commands. The cadp160 driver supports hot swap SCSI, hot plug PCI, 64-bit addressing (dual address cycle), domain validation, PCI bus clock rates up to 66MHZ and narrow and wide devices at 20MB/sec, 40MB/sec, 80MB/sec, and 160MB/sec.</p>				
FILES	<p>/platform/i86pc/kernel/drv/cadp160 32-bit ELF kernel module</p> <p>/platform/i86pc/kernel/drv/amd64/cadp160 64-bit ELF kernel module</p> <p>/boot/solaris/drivers/notisa.010/cadp160.bef realmode BEF driver</p> <p>/platform/i86pc/kernel/drv/cadp160.conf optional configuration file</p>				
ATTRIBUTES	<p>See attributes(5) for a description of the following attribute:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>x86</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	x86				
SEE ALSO	<p>prtconf(1M), driver.conf(4), pci(4), attributes(5), dlpi(7P), scsi_abort(9F), scsi_hba_attach(9F), scsi_ifgetcap(9F), scsi_ifsetcap(9F), scsi_reset(9F), scsi_sync_pkt(9F), scsi_transport(9F), scsi_device(9S), scsi_extended_sense(9S), scsi_inquiry(9S), scsi_pkt(9S)</p> <p><i>Writing Device Drivers</i></p> <p><i>Solaris (x86 Edition) Hardware Compatibility List</i></p> <p><i>ANSI Small Computer System Interface-2 (SCSI-2)</i></p>				

NAME	cadp – Adaptec Ultra-2 SCSI host bus adapter driver
SYNOPSIS	scsi@unit-address
DESCRIPTION	<p>The cadp host bus adapter driver is a SCSI-compliant nexus driver that supports the following Adaptec Ultra-2 SCSI devices:</p> <ul style="list-style-type: none"> ■ Adapters: Adaptec AHA-2940U2W, AHA-2940U2B, AHA-2940U2, AHA-2950U2B, AHA-3950U2B ■ Chips: AIC-7896/AIC-7897, AIC-7890/AIC-7890A, AIC-7891, AIC-7890AB, AIC-7890A <p>The cadp driver supports standard functions provided by the SCSI interface, including tagged and untagged queuing, Wide/Fast/Ultra SCSI, and auto request sense. The cadp driver does not support linked commands.</p>
PRECONFIGURE	<ul style="list-style-type: none"> ■ The Plug N Play SCAM Support option is not supported. ■ If the BIOS is enabled on the card, ensure that the Adaptec SCSISelect BIOS option Reset SCSI Bus at IC Initialization (under the Advanced Configuration Options menu) is set to Enabled. Run the SCSISelect utility by pressing Ctrl-A when you see the Adaptec banner during system boot. ■ If the adapter is being used in a multi-initiator configuration, do the following: (1) Ensure that the system boot disk is not on the shared (clustered) bus. (2) Set the Reset SCSI Bus at IC Initialization option to Disabled. (3) Set the Host Adapter BIOS option (under the Advanced Configuration Options menu) to Disabled:Not scan. (4) Add the allow-bus-reset=0 property to the /kernel/drv/cadp.conf file. ■ Reboot the system after you install patches.
Known Problems and Limitations	<ul style="list-style-type: none"> ■ The cadp.bef realmode driver supports only ten adapters at boot time. Therefore, ensure that the boot disk is attached to one of the first ten adapters. Note that all targets will be available for installation and use by the cadp driver. ■ Running the format(1M) command on a Seagate ST19171W 9 GB disk drive fails. ■ Some motherboards may have problems supporting channel B with boards based on the Adaptec AIC-7896 chip. The problem arises because the BIOS doesn't properly assign two interrupts for PCI interrupts INTA and INTB on the slot containing the AIC-7896 chip. As a result, timeouts and resets on those devices appear on the console. For some motherboards, you can work around the problem by setting the Advanced/PCI IRQ Mapping feature to ISA Legacy IRQs. ■ If you experience problems when using a narrow SCSI CD-ROM drive on the internal wide interface, disable "negotiate wide," "negotiate sync," or both for that device in the Adaptec configuration utility. ■ The Fujitsu narrow disk (M1603SAU) can reselect with an invalid queue tag ID. This violates the SCSI protocol and it causes the cadp driver to behave erroneously. Because this is difficult to guard against, you should disable tagged queuing for these targets. Use the iostat --E command to determine if you have a Fujitsu M1603S-512 disk. If you do, edit the /kernel/drv/cadp.conf file and

cadp(7D)

add the property `target n -scsi-options=0x1f78`, where n is the target number.

- The IBM external wide disk (DFHSS2W, Revision 1717) is not supported.
- When setting up a SCSI bus configuration, avoid connecting wide devices to a narrow bus. However, if you have such a configuration, add the following entry to the `cadp.conf` file: `target n -scsi-options=0x1df8` where n is the target ID of the wide device on the narrow bus. This entry disables wide negotiation for the specified target. Also ensure that the upper 8 bits of the bus are properly terminated at both ends of the SCSI chain.
- If you experience installation problems on systems with Intel 440BX/440GX motherboards, upgrade the motherboard BIOS with the latest revision.

CONFIGURATION

You configure the `cadp` host bus adapter driver by defining the properties found in `cadp.conf`. The `cadp.conf` file contains properties that you can modify, including: `scsi-options`, `target $\langle n \rangle$ -scsi-options`, `scsi-reset-delay`, and `scsi-initiator-id`. Properties in the `cadp.conf` file override global SCSI settings.

The property `target $\langle n \rangle$ -scsi-options` overrides the `scsi-options` property value for `target $\langle n \rangle$` , where $\langle n \rangle$ can vary from decimal 0 to 15. The `cadp` driver supports the following `scsi-options`: `SCSI_OPTIONS_DR`, `SCSI_OPTIONS_SYNC`, `SCSI_OPTIONS_TAG`, `SCSI_OPTIONS_FAST`, `SCSI_OPTIONS_WIDE`, `SCSI_OPTIONS_FAST20`, and `SCSI_OPTIONS_FAST40`.

You configure the SCSI devices using the Adaptec configuration utility. When configuring the devices, you should observe the following guidelines:

- Configure each device using a unique SCSI ID. On the Advanced Configuration Options menu, set Plug N Play SCAM Support to `Disabled`. Ensure that devices on either end of the SCSI chain are terminated. When mixing wide (16 bits) and narrow (8 bits) devices on the same wide chain, ensure that a wide device is at the end of the chain. If you place a narrow device at the end of the chain, wide devices on the same chain will terminate the low byte, resulting in an illegal configuration.
- If there is more than one controller, or an embedded controller, attempt to use one IRQ per controller.
- When prompted, enable bus mastering for the slot(s) with your host bus adapter(s.)
- Enable support for disks larger than 1 Gbyte, if applicable.

EXAMPLES

Create a file called `/kernel/drv/cadp.conf`, then add the following line:

```
scsi-options=0x78;
```

The above line disables tagged queuing, Fast/Ultra SCSI, and wide mode for all `cadp` instances.

To set `scsi-options` more specifically per target, add the following lines to `/kernel/drv/cadp.conf`:

```
target1-scsi-options=0x78;
device-type-scsi-options-list =
"SEAGATE ST32550W", "seagate-scsi-options" ;
seagate-scsi-options = 0x58;
scsi-options=0x3f8;
```

With the exception of one disk type that has `scsi-options` set to `0x58`, the above example sets `scsi-options` for target 1 to `0x78`, and all remaining targets to `0x3f8`.

The `scsi-options` properties that are specified per target ID have the highest precedence, followed by `scsi-options` per device type. Global `scsi-options` for all `cadp` instances per bus have the lowest precedence. You must reboot the system for the specified `scsi` options to take effect.

Driver Capabilities

To enable certain features on the `cadp` driver, the target driver must set capabilities. The following capabilities can be queried and modified by the target driver: `synchronous`, `tagged-qing`, `wide-xfer`, `auto-rqsense`, `qfull-retries`, and `qfull-retry-interval`. All other capabilities are query only.

By default, the `tagged-qing`, `auto-rqsense`, and `wide-xfer` capabilities are disabled. The `disconnect`, `synchronous`, and `untagged-qing` capabilities are always enabled. The `cadp` driver capabilities can only be assigned binary values (0 or 1). The default value for `qfull-retries` is `10` and the default value for `qfull-retry-interval` is `100`. The `qfull-retries` capability is `au_char` (0 to 255) while `qfull-retry-interval` is a `u_short` (0 to 65535).

If a conflict occurs between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only `whom != 0` is supported in the `scsi_ifsetcap(9F)` call. See `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

FILES /kernel/drv/cadp ELF kernel module
/kernel/drv/cadp.conf Optional configuration file

ATTRIBUTES See `attributes(5)` for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `prtconf(1M)`, `driver.conf(4)`, `pci(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

Solaris (Intel Platform Edition) Hardware Compatibility List

ANSI Small Computer System Interface-2 (SCSI-2)

cadp(7D)

NOTES | The cadp driver supports the adapters and chipsets listed in this man page. For information on support of additional devices, see the *Solaris (Intel Platform Edition) Hardware Compatibility List* a component of the *Information Library for Solaris 8 (Intel Platform Edition)*.

The cadp driver exports properties indicating (per target) the negotiated transfer speed (`target<n>-sync-speed`), whether wide bus (`target<n>-wide`), is supported for that particular target (`target<n>-scsi-options`), and whether tagged queuing (`target<n>-tag-queue`) has been enabled. The `sync-speed` property value is the data transfer rate in KB/sec. The `target<n>-tag-queue` and the `target<n>-wide` property have value 1 to indicate that the corresponding capability is enabled, or 0 to indicate that the capability is disabled. See `prtconf(1M)` (verbose option) for information on viewing the cadp properties.

Sample output is provided below:

```
pci9005,f500, instance #2
System software properties:
  name <interrupt-priorities> length <4>
    value <0x05000000>.
  name <tape> length <5>
    value <0x7363747000>.
  name <disk> length <5>
    value <0x7363646b00>.
  name <queue> length <6>
    value <0x71736f727400>.
  name <flow_control> length <6>
    value <0x646d756c7400>.
Driver properties:
  name <target0-tag-queue> length <4>
    value <0x01000000>.
  name <target0-wide> length <4>
    value <0x01000000>.
  name <target0-sync-speed> length <4>
    value <0x28000000>.
  name <chosen-interrupt> length <8>
    value <0x0100000000000000>.
  name <scsi-selection-timeout> length <4>
    value <0xfa000000>.
  name <scsi-options> length <4>
    value <0xf81f0000>.
  name <scsi-watchdog-tick> length <4>
    value <0x0a000000>.
  name <scsi-tag-age-limit> length <4>
    value <0x02000000>.
  name <scsi-reset-delay> length <4>
    value <0xb80b0000>.
```


NAME	cdio – CD-ROM control operations
SYNOPSIS	<code>#include <sys/cdio.h></code>
DESCRIPTION	<p>The set of <code>ioctl(2)</code> commands described below are used to perform audio and CD-ROM specific operations. Basic to these <code>cdio</code> <code>ioctl</code> requests are the definitions in <code><sys/cdio.h></code>.</p> <p>Several CD-ROM specific commands can report addresses either in LBA (Logical Block Address) format or in MSF (Minute, Second, Frame) format. The <code>READ HEADER</code>, <code>READ SUBCHANNEL</code>, and <code>READ TABLE OF CONTENTS</code> commands have this feature.</p> <p>LBA format represents the logical block address for the CD-ROM absolute address field or for the offset from the beginning of the current track expressed as a number of logical blocks in a CD-ROM track relative address field. MSF format represents the physical address written on CD-ROM discs, expressed as a sector count relative to either the beginning of the medium or the beginning of the current track.</p>
IOCTLS	<p>The following I/O controls do not have any additional data passed into or received from them.</p> <p><code>CDROMSTART</code> This <code>ioctl()</code> spins up the disc and seeks to the last address requested.</p> <p><code>CDROMSTOP</code> This <code>ioctl()</code> spins down the disc.</p> <p><code>CDROMPAUSE</code> This <code>ioctl()</code> pauses the current audio play operation.</p> <p><code>CDROMRESUME</code> This <code>ioctl()</code> resumes the paused audio play operation.</p> <p><code>CDROMEJECT</code> This <code>ioctl()</code> ejects the caddy with the disc.</p> <p><code>CDROMCLOSETRAY</code> This <code>ioctl()</code> closes the caddy with the disc.</p> <p>The following I/O controls require a pointer to the structure for that <code>ioctl()</code>, with data being passed into the <code>ioctl()</code>.</p> <p><code>CDROMPLAYMSF</code> This <code>ioctl()</code> command requests the drive to output the audio signals at the specified starting address and continue the audio play until the specified ending address is detected. The address is in MSF format. The third argument of this <code>ioctl()</code> call is a pointer to the type <code>struct cdrom_msf</code>.</p> <pre> /* * definition of play audio msf structure */ struct cdrom_msf { unsigned char cdmsf_min0; /* starting minute*/ unsigned char cdmsf_sec0; /* starting second*/ </pre>

cdio(7I)

```
    unsigned char    cdmsf_frame0;    /*starting frame*/
    unsigned char    cdmsf_min1;     /* ending minute */
    unsigned char    cdmsf_sec1;     /* ending second */
    unsigned char    cdmsf_frame1;    /* ending frame */
};
```

The CDROMREADTOCENTRY ioctl request may be used to obtain the start time for a track. An approximation of the finish time can be obtained by using the CDROMREADTOCENTRY ioctl request to retrieve the start time of the track following the current track.

The leadout track is the next consecutive track after the last audio track. Hence, the start time of the leadout track may be used as the effective finish time of the last audio track.

CDROMPLAYTRKIND

This ioctl() command is similar to CDROMPLAYMSF. The starting and ending address is in track/index format. The third argument of the ioctl() call is a pointer to the type struct cdrom_ti.

```
/*
 * definition of play audio track/index structure
 */
struct cdrom_ti {
    unsigned char    cdti_trk0;      /* starting track*/
    unsigned char    cdti_ind0;      /* starting index*/
    unsigned char    cdti_trk1;      /* ending track */
    unsigned char    cdti_ind1;      /* ending index */
};
```

CDROMVOLCTRL

This ioctl() command controls the audio output level. The SCSI command allows the control of up to four channels. The current implementation of the supported CD-ROM drive only uses channel 0 and channel 1. The valid values of volume control are between 0x00 and 0xFF, with a value of 0xFF indicating maximum volume. The third argument of the ioctl() call is a pointer to struct cdrom_volctrl which contains the output volume values.

```
/*
 * definition of audio volume control structure
 */
struct cdrom_volctrl {
    unsigned char    channel0;
    unsigned char    channel1;
    unsigned char    channel2;
    unsigned char    channel3;
};
```

The following I/O controls take a pointer that will have data returned to the user program from the CD-ROM driver.

CDROMREADTOCHDR

This `ioctl()` command returns the header of the table of contents (TOC). The header consists of the starting tracking number and the ending track number of the disc. These two numbers are returned through a pointer of `struct cdrom_tochdr`. While the disc can start at any number, all tracks between the first and last tracks are in contiguous ascending order.

```
/*
 * definition of read toc header structure
 */
struct cdrom_tochdr {
    unsigned char    cdth_trk0;    /* starting track*/
    unsigned char    cdth_trk1;    /* ending track*/
};
```

CDROMREADTOCENTRY

This `ioctl()` command returns the information of a specified track. The third argument of the function call is a pointer to the type `struct cdrom_tocentry`. The caller needs to supply the track number and the address format. This command will return a 4-bit `adr` field, a 4-bit `ctrl` field, the starting address in MSF format or LBA format, and the data mode if the track is a data track. The `ctrl` field specifies whether the track is data or audio.

```
/*
 * definition of read toc entry structure
 */
struct cdrom_tocentry {
    unsigned char    cdte_track;
    unsigned char    cdte_adr    :4;
    unsigned char    cdte_ctrl   :4;
    unsigned char    cdte_format;
    union {
        struct {
            unsigned char    minute;
            unsigned char    second;
            unsigned char    frame;
        } msf;
        int    lba;
    } cdte_addr;
    unsigned char    cdte_datamode;
};
```

To get the information from the leadout track, the following value is appropriate for the `cdte_track` field:

```
CDROM_LEADOUT          Leadout track
```

To get the information from the data track, the following value is appropriate for the `cdte_ctrl` field:

```
CDROM_DATA_TRACK      Data track
```

The following values are appropriate for the `cdte_format` field:

```
CDROM_LBA              LBA format
```

CDROM_MSF MSF format

CDROMSUBCHNL

This `ioctl()` command reads the Q sub-channel data of the current block. The subchannel data includes track number, index number, absolute CD-ROM address, track relative CD-ROM address, control data and audio status. All information is returned through a pointer to `struct cdrom_subchnl`. The caller needs to supply the address format for the returned address.

```
struct cdrom_subchnl {
    unsigned char    cdsc_format;
    unsigned char    cdsc_audiostatus;
    unsigned char    cdsc_adr:    4;
    unsigned char    cdsc_ctrl:   4;
    unsigned char    cdsc_trk;
    unsigned char    cdsc_ind;
    union {
        struct {
            unsigned char    minute;
            unsigned char    second;
            unsigned char    frame;
        } msf;
        int    lba;
    } cdsc_absaddr;
    union {
        struct {
            unsigned char    minute;
            unsigned char    second;
            unsigned char    frame;
        } msf;
        int    lba;
    } cdsc_reladdr;
};
```

The following values are valid for the audio status field returned from `READ SUBCHANNEL` command:

CDROM_AUDIO_INVALID	Audio status not supported.
CDROM_AUDIO_PLAY	Audio play operation in progress.
CDROM_AUDIO_PAUSED	Audio play operation paused.
CDROM_AUDIO_COMPLETED	Audio play successfully completed.
CDROM_AUDIO_ERROR	Audio play stopped due to error.
CDROM_AUDIO_NO_STATUS	No current audio status to return.

CDROMREADOFFSET

This `ioctl()` command returns the absolute CD-ROM address of the first track in the last session of a Multi-Session CD-ROM. The third argument of the `ioctl()` call is a pointer to an `int`.

CDROMCDDA

This `ioctl()` command returns the CD-DA data or the subcode data. The third argument of the `ioctl()` call is a pointer to the type `struct cdrom_cdda`. In

addition to allocating memory and supplying its address, the caller needs to supply the starting address of the data, the transfer length in terms of the number of blocks to be transferred, and the subcode options. The caller also needs to issue the `CDROMREADTOCENTRY ioctl()` to find out which tracks contain CD-DA data before issuing this `ioctl()`.

```
/*
 * Definition of CD-DA structure
 */
struct cdrom_cdda {
    unsigned int    cdda_addr;
    unsigned int    cdda_length;
    caddr_t         cdda_data;
    unsigned char   cdda_subcode;
};
```

`cdda_addr` signifies the starting logical block address.

`cdda_length` signifies the transfer length in blocks. The length of the block depends on the `cdda_subcode` selection, which is explained below.

To get the subcode information related to CD-DA data, the following values are appropriate for the `cdda_subcode` field:

<code>CDROM_DA_NO_SUBCODE</code>	CD-DA data with no subcode.
<code>CDROM_DA_SUBQ</code>	CD-DA data with sub Q code.
<code>CDROM_DA_ALL_SUBCODE</code>	CD-DA data with all subcode.
<code>CDROM_DA_SUBCODE_ONLY</code>	All subcode only.

To allocate the memory related to CD-DA and/or subcode data, the following values are appropriate for each data block transferred:

CD-DA data with no subcode	2352 bytes
CD-DA data with sub Q code	2368 bytes
CD-DA data with all subcode	2448 bytes
All subcode only	96 bytes

CDROMCDXA

This `ioctl()` command returns the CD-ROM XA (CD-ROM Extended Architecture) data according to CD-ROM XA format. The third argument of the `ioctl()` call is a pointer to the type `struct cdrom_cdxa`. In addition to allocating memory and supplying its address, the caller needs to supply the starting address of the data, the transfer length in terms of number of blocks, and the format. The caller also needs to issue the `CDROMREADTOCENTRY ioctl()` to find out which tracks contain CD-ROM XA data before issuing this `ioctl()`.

```
/*
 * Definition of CD-ROM XA structure
 */
```

cdio(7I)

```
struct cdrom_cdxa {
    unsigned int    cdx_a_addr;
    unsigned int    cdx_a_length;
    caddr_t         cdx_a_data;
    unsigned char   cdx_a_format;
};
```

To get the proper CD-ROM XA data, the following values are appropriate for the `cdxa_format` field:

CDROM_XA_DATA	CD-ROM XA data only
CDROM_XA_SECTOR_DATA	CD-ROM XA all sector data
CDROM_XA_DATA_W_ERROR	CD-ROM XA data with error flags data

To allocate the memory related to CD-ROM XA format, the following values are appropriate for each data block transferred:

CD-ROM XA data only	2048 bytes
CD-ROM XA all sector data	2352 bytes
CD-ROM XA data with error flags data	2646 bytes

CDROMSUBCODE

This `ioctl()` command returns raw subcode data (subcodes P ~ W are described in the "Red Book," see SEE ALSO) to the initiator while the target is playing audio. The third argument of the `ioctl()` call is a pointer to the type `struct cdrom_subcode`. The caller needs to supply the transfer length in terms of number of blocks and allocate memory for subcode data. The memory allocated should be a multiple of 96 bytes depending on the transfer length.

```
/*
 * Definition of subcode structure
 */
struct cdrom_subcode {
    unsigned int    cdsc_length;
    caddr_t         cdsc_addr;
};
```

The next group of I/O controls get and set various CD-ROM drive parameters.

CDROMGBLKMDE

This `ioctl()` command returns the current block size used by the CD-ROM drive. The third argument of the `ioctl()` call is a pointer to an integer.

CDROMSBLKMDE

This `ioctl()` command requests the CD-ROM drive to change from the current block size to the requested block size. The third argument of the `ioctl()` call is an integer which contains the requested block size.

This `ioctl()` command operates in exclusive-use mode only. The caller must ensure that no other processes can operate on the same CD-ROM device before issuing this `ioctl().read(2)` behavior subsequent to this `ioctl()` remains the same: the caller is still constrained to read the raw device on block boundaries and in block multiples.

To set the proper block size, the following values are appropriate:

<code>CDROM_BLK_512</code>	512 bytes
<code>CDROM_BLK_1024</code>	1024 bytes
<code>CDROM_BLK_2048</code>	2048 bytes
<code>CDROM_BLK_2056</code>	2056 bytes
<code>CDROM_BLK_2336</code>	2336 bytes
<code>CDROM_BLK_2340</code>	2340 bytes
<code>CDROM_BLK_2352</code>	2352 bytes
<code>CDROM_BLK_2368</code>	2368 bytes
<code>CDROM_BLK_2448</code>	2448 bytes
<code>CDROM_BLK_2646</code>	2646 bytes
<code>CDROM_BLK_2647</code>	2647 bytes

`CDROMGDRVSPEED`

This `ioctl()` command returns the current CD-ROM drive speed. The third argument of the `ioctl()` call is a pointer to an integer.

`CDROMSDRVSPEED`

This `ioctl()` command requests the CD-ROM drive to change the current drive speed to the requested drive speed. This speed setting is only applicable when reading data areas. The third argument of the `ioctl()` is an integer which contains the requested drive speed.

To set the CD-ROM drive to the proper speed, the following values are appropriate:

<code>CDROM_NORMAL_SPEED</code>	150k/second
<code>CDROM_DOUBLE_SPEED</code>	300k/second
<code>CDROM_QUAD_SPEED</code>	600k/second
<code>CDROM_MAXIMUM_SPEED</code>	300k/second (2x drive) 600k/second (4x drive)

Note that these numbers are only accurate when reading 2048 byte blocks. The CD-ROM drive will automatically switch to normal speed when playing audio tracks and will switch back to the speed setting when accessing data.

SEE ALSO `ioctl(2)`, `read(2)`

cdio(7I)

N. V. Phillips and Sony Corporation, *System Description Compact Disc Digital Audio*, ("Red Book").

N. V. Phillips and Sony Corporation, *System Description of Compact Disc Read Only Memory*, ("Yellow Book").

N. V. Phillips, Microsoft, and Sony Corporation, *System Description CD-ROM XA*, 1991.

Volume and File Structure of CD-ROM for Information Interchange, ISO 9660:1988(E).

SCSI-2 Standard, document X3T9.2/86-109

SCSI Multimedia Commands, Version 2 (MMC-2)

NOTES

The CDROMCDDA, CDROMCDXA, CDROMSUBCODE, CDROMGDRVSPEED, CDROMSDRVSPEED, and some of the block sizes in CDROMSBLKMODE are designed for new Sun-supported CD-ROM drives and might not work on some of the older CD-ROM drives.

CDROMCDDA, CDROMCDXA and CDROMSUBCODE will return error if the transfer length exceeds valid limits as determined appropriate. Example: for MMC-2 drives, length can not exceed 3 bytes (i.e. 0xfffff). The same restriction is enforced for older, pre-MMC-2 drives, as no limit was published for these older drives (and 3 bytes is reasonable for all media). Note that enforcing this limit does not imply that values passed in below this limit will actually be applicable for each and every piece of media.

The interface to this device is preliminary and subject to change in future releases. Programs should be written in a modular fashion so that future changes can be easily incorporated.

NAME	ce – Cassini Gigabit-Ethernet device driver
SYNOPSIS	/dev/ce
DESCRIPTION	<p>The <code>ce</code> Sun Gigabit-Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code>, over all implementations of PCI Cassini Gigabit-Ethernet add-in adapters. Multiple Cassini-based adapters installed within the system are supported by the driver. The <code>ce</code> driver provides basic support for the Cassini-based Ethernet hardware and handles the <code>pci108e,abba</code> (PCI Cassini) devices. Functions include chip initialization, frame transmit and receive, multicast and promiscuous support, and error recovery and reporting. The Cassini device provides 1000BASE-SX networking interfaces using the Cassini ASIC external SERDES and fiber optical transceiver, or 10/100/1000BASE-T using a Cassini ASIC attached to a GMII twisted pair copper transceiver, or 10/100BASE-T using a Cassini ASIC attached to a MII twisted pair copper transceiver.</p> <p>The 1000Base-SX standard specifies an auto-negotiation protocol to automatically select the mode of operation. In addition to the duplex mode of operation, the Cassini ASIC can auto-negotiate for <i>IEEE 802.3x</i> frame-based flow control capabilities. The Cassini PCS can perform auto-negotiation with the link's remote-end (link partner) and receives the capabilities of the remote end. It selects the highest common denominator mode of operation based on the priorities. It also supports forced-mode of operation where the driver selects the mode of operation.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The <code>/dev/ce</code> cloning character-special device is used to access all <code>ce</code> controllers installed on the system.</p>
ce and DLPI	<p>The <code>ce</code> driver is a Style 2 data link service provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code> type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <code><sys/dlpi.h></code>. Refer to <code>dlpi(7P)</code> for more information. An explicit <code>DL_ATTACH_REQ</code> message by the user is required to associate the opened stream with a particular device (<code>ppa</code>). The <code>ppa</code> ID is interpreted as an <code>unsigned long</code> data type and indicates the corresponding device instance (unit) number. An error (<code>DL_ERROR_ACK</code>) is returned by the driver if the <code>ppa</code> field value does not correspond to a valid device instance number for this system. The device is initialized on first attach and de-initialized (stopped) upon last detach.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> are as follows:</p> <ul style="list-style-type: none"> ■ Maximum SDU is 1500 (<code>ETHERMTU</code> - defined in <code><sys/ethernet.h></code>). ■ Minimum SDU is 0. ■ The <code>dlsap</code> address length is 8. ■ MAC type is <code>DL_ETHER</code>. ■ The <code>sap</code> length value is -2 meaning the physical address component is followed immediately by a 2 byte <code>sap</code> component within the DLSAP address.

ce(7D)

- Service mode is DL_CLDLS.
- Optional quality of service (QOS) is not supported; the QOS fields are 0.
- Provider style is DL_STYLE2.
- Version is DL_VERSION_2.
- Broadcast address value is Ethernet/IEEE broadcast address (0xFFFFF).

Once in the DL_ATTACHED state, you must send a DL_BIND_REQ to associate a particular SAP (Service Access Pointer) with the stream. The ce driver interprets the sap field within the DL_BIND_REQ as an Ethernet "type," therefore valid values for the sap field are in the range [0-0xFFFF]. Only one Ethernet type can be bound to the stream at any time.

If you select a sap with a value of 0, the receiver will be in 802.3 mode. All frames received from the media having a "type" field in the range [0-1500] are assumed to be 802.3 frames and are routed up all open streams which are bound to sap value 0. If more than one stream is in 802.3 mode, the frame will be duplicated and routed up multiple streams as DL_UNITDATA_IND messages.

In transmission, the driver checks the sap field of the DL_BIND_REQ to verify that the sap value is 0, and that the destination type field is in the range [0-1500]. If either is true, the driver computes the length of the message, not including initial M_PROTO mblk (message block), of all subsequent DL_UNITDATA_REQ messages and transmits 802.3 frames that have this value in the MAC frame header length field.

The ce driver DLSAP address format consists of the 6 byte physical (Ethernet) address component followed immediately by the 2 byte sap (type) component producing an 8 byte DLSAP address. Applications should *not* hard code to this particular implementation-specific DLSAP address format, but use information returned in the DL_INFO_ACK primitive to compose and decompose DLSAP addresses. The sap length, full DLSAP length, and sap/physical ordering are included within the DL_INFO_ACK. The physical address length can be computed by subtracting the sap length from the full DLSAP address length or by issuing the DL_PHYS_ADDR_REQ to obtain the current physical address associated with the stream.

Once in the DL_BOUND state, you can transmit frames on the Ethernet by sending DL_UNITDATA_REQ messages to the ce driver. The ce driver will route received Ethernet frames up all open and bound streams having a sap which matches the Ethernet type as DL_UNITDATA_IND messages. Received Ethernet frames are duplicated and routed up multiple open streams, if necessary. The DLSAP address contained within the DL_UNITDATA_REQ and DL_UNITDATA_IND messages consists of both the sap (type) and physical (Ethernet) components.

In addition to the mandatory connectionless DLPI message set, the driver additionally supports the following primitives.

ce Primitives	<p>The <code>DL_ENABMULTI_REQ</code> and <code>DL_DISABMULTI_REQ</code> primitives enable/disable reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. These primitives are accepted by the driver in any state following <code>DL_ATTACHED</code>.</p> <p>The <code>DL_PROMISCON_REQ</code> and <code>DL_PROMISCOFF_REQ</code> primitives with the <code>DL_PROMISC_PHYS</code> option set in the <code>dl_level</code> field enables/disables reception of all “promiscuous mode” frames on the media, including frames generated by the local host. When used with the <code>DL_PROMISC_SAP</code> option set, this enables/disables reception of all <code>sap</code> (Ethernet type) values. When used with the <code>DL_PROMISC_MULTI</code> option set this enables/disables reception of all multicast group addresses. The effect of each is always on a per-stream basis and independent of the other <code>sap</code> and physical level configurations on this stream or other streams.</p> <p>The <code>DL_PHYS_ADDR_REQ</code> primitive returns the 6 octet Ethernet address currently associated (attached) to the stream in the <code>DL_PHYS_ADDR_ACK</code> primitive. This primitive is valid only in states following a successful <code>DL_ATTACH_REQ</code>.</p> <p>The <code>DL_SET_PHYS_ADDR_REQ</code> primitive changes the 6 octet Ethernet address currently associated (attached) to this stream. The credentials of the process which originally opened this stream must be superuser. Otherwise <code>EPERM</code> is returned in the <code>DL_ERROR_ACK</code>. This primitive is destructive because it affects all other current and future streams attached to this device. An <code>M_ERROR</code> is sent up all other streams attached to this device when this primitive is successful on this stream. Once changed, all streams subsequently opened and attached to this device will obtain this new physical address. Once changed, the physical address will remain until this primitive is used to change the physical address again or the system is rebooted, whichever comes first.</p>
ce DRIVER	<p>By default, the <code>ce</code> driver performs auto-negotiation to select the mode and flow control capabilities of the link.</p> <p>The link can assume one of the following modes:</p> <ul style="list-style-type: none"> ■ 1000 Mbps, full-duplex ■ 1000 Mbps, half-duplex ■ Symmetric pause ■ Asymmetric pause <p>Speeds and modes are described in the 1000Base-TX standard.</p> <p>The auto-negotiation protocol automatically selects:</p> <ul style="list-style-type: none"> ■ Operation mode (half-duplex or full-duplex) ■ Flow control capability (symmetric and/or asymmetric) <p>The auto-negotiation protocol does the following:</p> <ul style="list-style-type: none"> ■ Gets all modes of operation supported by the link partner. ■ Advertises its capabilities to the link partner.

ce(7D)

- Selects the highest common denominator mode of operation based on the priorities.

The Cassini hardware can operate in all modes listed above, providing auto-negotiation is used by default to bring up the link and select the common mode of operation with the link partner. The PCS also supports forced-mode of operation in which the driver can select the mode of operation and the flow control capabilities, using the `ndd` utility.

The Cassini device also supports programmable IPG (Inter-Packet Gap) parameters `ipg1` and `ipg2`. By default, the driver sets `ipg1` and `ipg2` to 8 and 4 byte-times respectively (which are the standard values). If desired, you can alter these values from the standard 1000 Mbps IPG set to 0.096 microseconds.

ce Parameter List

The `ce` driver enables the setting and getting of various parameters for the Cassini device. The parameter list includes *current transceiver status*, *current link status*, *inter-packet gap*, *PCS capabilities* and *link partner capabilities*.

The PCS features two set of capabilities. One set reflects the capabilities of the hardware and are read-only. The second set, which reflects the values you choose, are used in speed selection and possess read/write capabilities. At boot time, these two sets of capabilities are the same. The link partner capabilities are also read-only because the current default value of these parameters can be read but not modified.

FILES

`/dev/ce`
ce special character device.

`/platform/sun4u/kernel/drv/ce.conf`
SPARC system-wide default device driver properties.

`/platform/amd64/kernel/drv/ce.conf`
64-bit x86 system-wide default device driver properties.

`/kernel/drv/ce.conf`
32-bit x86 system-wide default device driver properties.

`/kernel/drv/amd64/ce.conf`
64-bit x86 system-wide default device driver properties.

SEE ALSO

`ndd(1M)`, `netstat(1M)`, `driver.conf(4)`, [ge\(7D\)](#), [hme\(7D\)](#), [qfe\(7D\)](#), [dlpi\(7P\)](#)

NAME	cgsix – accelerated 8-bit color frame buffer
SYNOPSIS	<code>/dev/fbs/cgsixn</code>
DESCRIPTION	<p>cgsix is a low-end graphics accelerator designed to enhance vector and polygon drawing performance. It has an 8-bit color frame buffer and provides the standard frame buffer interface as defined in <code>fbio(7I)</code>.</p> <p>In addition, cgsix supports the following cgsix-specific IOCTL, defined in <code><sys/fbio.h></code>.</p> <p><code>FBIOGXINFO</code> Returns cgsix-specific information about the hardware. See the definition of <code>cg6_info</code> in <code><sys/fbio.h></code> for more information.</p> <p>cgsix has registers and memory that may be mapped with <code>mmap(2)</code>, using the offsets defined in <code><sys/cg6reg.h></code>.</p>
FILES	<code>/dev/fbs/cgsix0</code>
SEE ALSO	<code>mmap(2)</code> , <code>fbio(7I)</code>

chs(7D)

NAME	chs – IBM ServeRAID PCI host adapter driver														
DESCRIPTION	The chs driver is the IBM ServeRAID PCI controller driver.														
PRECONFIGURATION															
Known Problems and Limitations	To prevent data loss, a SCSI disk drive that is not defined as part of any physical pack within a logical drive will not be accessible through the Solaris operating environment.														
CONFIGURATION	<p>Auto-configuration code determines whether the adapter is present at the configured address and what types of devices are attached to it. The IBM ServeRAID is primarily used as a disk array (system drive) controller.</p> <p>To configure the attached disk arrays, you must configure the controller (using the configuration utilities provided by the hardware manufacturer) before booting the Solaris operating environment. You use the configuration utilities to set RAID levels, stripe parameters, cache mechanisms and perform other functions. For more information, see the user manual supplied with your hardware.</p>														
FILES	<table><tr><td>/kernel/drv/chs.conf</td><td>chs configuration file</td></tr><tr><td>/dev/dsk/cndn [s p] n</td><td>block device</td></tr><tr><td>/dev/rdisk/cndn [s p] n</td><td>raw device where:</td></tr><tr><td>cn</td><td>controller n</td></tr><tr><td>dn</td><td>LUN n (0–7)</td></tr><tr><td>sn</td><td>UNIX system slice n (0–15)</td></tr><tr><td>pn</td><td>fdisk(1M) partition (0)</td></tr></table>	/kernel/drv/chs.conf	chs configuration file	/dev/dsk/cndn [s p] n	block device	/dev/rdisk/cndn [s p] n	raw device where:	cn	controller n	dn	LUN n (0–7)	sn	UNIX system slice n (0–15)	pn	fdisk(1M) partition (0)
/kernel/drv/chs.conf	chs configuration file														
/dev/dsk/cndn [s p] n	block device														
/dev/rdisk/cndn [s p] n	raw device where:														
cn	controller n														
dn	LUN n (0–7)														
sn	UNIX system slice n (0–15)														
pn	fdisk(1M) partition (0)														
ATTRIBUTES	See attributes(5) for a description of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>x86</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86										
ATTRIBUTE TYPE	ATTRIBUTE VALUE														
Architecture	x86														
SEE ALSO	fdisk(1M), attributes(5), cmdk(7D)														

NAME	cmdk – common disk driver																		
SYNOPSIS	<code>cmdk@target, lun : [partition slice]</code>																		
DESCRIPTION	<p>The <code>cmdk</code> device driver is a common interface to various disk devices. The driver supports magnetic fixed disks and magnetic removable disks.</p> <p>The block-files access the disk using the system's normal buffering mechanism and are read and written without regard to physical disk records. There is also a "raw" interface that provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in one I/O operation; raw I/O is therefore considerably more efficient when many bytes are transmitted. The names of the block files are found in <code>/dev/dsk</code>; the names of the raw files are found in <code>/dev/rdisk</code>.</p> <p>I/O requests to the magnetic disk must have an offset and transfer length that is a multiple of 512 bytes or the driver returns an <code>EINVAL</code> error.</p> <p>Slice 0 is normally used for the root file system on a disk, slice 1 as a paging area (for example, swap), and slice 2 for backing up the entire fdisk partition for Solaris software. Other slices may be used for <code>usr</code> file systems or system reserved area.</p> <p>Fdisk partition 0 is to access the entire disk and is generally used by the <code>fdisk(1M)</code> program.</p>																		
FILES	<table border="0"> <tr> <td><code>/dev/dsk/cndn[s p]n</code></td> <td>block device (IDE)</td> </tr> <tr> <td><code>/dev/rdisk/cndn[s p]n</code></td> <td>raw device (IDE)</td> </tr> <tr> <td></td> <td>where:</td> </tr> <tr> <td><code>cn</code></td> <td>controller <i>n</i></td> </tr> <tr> <td><code>dn</code></td> <td>lun <i>n</i> (0-7)</td> </tr> <tr> <td><code>sn</code></td> <td>UNIX system slice <i>n</i> (0-15)</td> </tr> <tr> <td><code>pn</code></td> <td>fdisk partition (0)</td> </tr> <tr> <td><code>/kernel/drv/cmdk</code></td> <td>32-bit kernel module.</td> </tr> <tr> <td><code>/kernel/drv/amd64/cmdk</code></td> <td>64-bit kernel module.</td> </tr> </table>	<code>/dev/dsk/cndn[s p]n</code>	block device (IDE)	<code>/dev/rdisk/cndn[s p]n</code>	raw device (IDE)		where:	<code>cn</code>	controller <i>n</i>	<code>dn</code>	lun <i>n</i> (0-7)	<code>sn</code>	UNIX system slice <i>n</i> (0-15)	<code>pn</code>	fdisk partition (0)	<code>/kernel/drv/cmdk</code>	32-bit kernel module.	<code>/kernel/drv/amd64/cmdk</code>	64-bit kernel module.
<code>/dev/dsk/cndn[s p]n</code>	block device (IDE)																		
<code>/dev/rdisk/cndn[s p]n</code>	raw device (IDE)																		
	where:																		
<code>cn</code>	controller <i>n</i>																		
<code>dn</code>	lun <i>n</i> (0-7)																		
<code>sn</code>	UNIX system slice <i>n</i> (0-15)																		
<code>pn</code>	fdisk partition (0)																		
<code>/kernel/drv/cmdk</code>	32-bit kernel module.																		
<code>/kernel/drv/amd64/cmdk</code>	64-bit kernel module.																		
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:																		
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>x86</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	x86														
ATTRIBUTE TYPE	ATTRIBUTE VALUE																		
Architecture	x86																		
SEE ALSO	<code>fdisk(1M)</code> , <code>mount(1M)</code> , <code>lseek(2)</code> , <code>read(2)</code> , <code>write(2)</code> , <code>readdir(3C)</code> , <code>scsi(4)</code> , <code>vfstab(4)</code> , <code>attributes(5)</code> , <code>dkio(7I)</code>																		

connld(7M)

NAME	connld – line discipline for unique stream connections														
SYNOPSIS	<pre>#include </sys/steam.h> int ioctl (fd, I_PUSH, "connld");</pre>														
DESCRIPTION	<p>connld is a STREAMS-based module that provides unique connections between server and client processes. It can only be pushed (see streamio(7I)) onto one end of a STREAMS-based pipe that may subsequently be attached to a name in the file system name space with fattach(3C). After the pipe end is attached, a new pipe is created internally when an originating process attempts to open(2) or creat(2) the file system name. A file descriptor for one end of the new pipe is packaged into a message identical to that for the ioctl I_SENDFD (see streamio(7I)) and is transmitted along the stream to the server process on the other end. The originating process is blocked until the server responds.</p> <p>The server responds to the I_SENDFD request by accepting the file descriptor through the I_RECVFD ioctl message. When this happens, the file descriptor associated with the other end of the new pipe is transmitted to the originating process as the file descriptor returned from open(2) or creat(2).</p> <p>If the server does not respond to the I_SENDFD request, the stream that the connld module is pushed on becomes uni-directional because the server will not be able to retrieve any data off the stream until the I_RECVFD request is issued. If the server process exits before issuing the I_RECVFD request, the open(2) or the creat(2) invocation will fail and return -1 to the originating process.</p> <p>When the connld module is pushed onto a pipe, it ignores messages going back and forth through the pipe.</p>														
ERRORS	<p>On success, an open of connld returns 0. On failure, errno is set to the following values:</p> <table><tr><td>EINVAL</td><td>A stream onto which connld is being pushed is not a pipe or the pipe does not have a write queue pointer pointing to a stream head read queue.</td></tr><tr><td>EINVAL</td><td>The other end of the pipe onto which connld is being pushed is linked under a multiplexor.</td></tr><tr><td>EPIPE</td><td>connld is being pushed onto a pipe end whose other end is no longer there.</td></tr><tr><td>ENOMEM</td><td>An internal pipe could not be created.</td></tr><tr><td>ENXIO</td><td>An M_HANGUP message is at the stream head of the pipe onto which connld is being pushed.</td></tr><tr><td>EAGAIN</td><td>Internal data structures could not be allocated.</td></tr><tr><td>ENFILE</td><td>A file table entry could not be allocated.</td></tr></table>	EINVAL	A stream onto which connld is being pushed is not a pipe or the pipe does not have a write queue pointer pointing to a stream head read queue.	EINVAL	The other end of the pipe onto which connld is being pushed is linked under a multiplexor.	EPIPE	connld is being pushed onto a pipe end whose other end is no longer there.	ENOMEM	An internal pipe could not be created.	ENXIO	An M_HANGUP message is at the stream head of the pipe onto which connld is being pushed.	EAGAIN	Internal data structures could not be allocated.	ENFILE	A file table entry could not be allocated.
EINVAL	A stream onto which connld is being pushed is not a pipe or the pipe does not have a write queue pointer pointing to a stream head read queue.														
EINVAL	The other end of the pipe onto which connld is being pushed is linked under a multiplexor.														
EPIPE	connld is being pushed onto a pipe end whose other end is no longer there.														
ENOMEM	An internal pipe could not be created.														
ENXIO	An M_HANGUP message is at the stream head of the pipe onto which connld is being pushed.														
EAGAIN	Internal data structures could not be allocated.														
ENFILE	A file table entry could not be allocated.														
SEE ALSO	creat(2) , open(2) , fattach(3C) , streamio(7I)														

console(7D)

NAME	console – STREAMS-based console interface
SYNOPSIS	<code>/dev/console</code>
DESCRIPTION	<p>The file <code>/dev/console</code> refers to the system console device. <code>/dev/console</code> should be used for interactive purposes only. Use of <code>/dev/console</code> for logging purposes is discouraged; <code>syslog(3C)</code> or <code>msglog(7D)</code> should be used instead.</p> <p>The identity of this device depends on the EEPROM or NVRAM settings in effect at the most recent system reboot; by default, it is the “workstation console” device consisting of the workstation keyboard and frame buffer acting in concert to emulate an ASCII terminal (see <code>wscons(7D)</code>).</p> <p>Regardless of the system configuration, the console device provides asynchronous serial driver semantics so that, in conjunction with the STREAMS line discipline module <code>ldterm(7M)</code>, it supports the <code>termio(7I)</code> terminal interface.</p>
SEE ALSO	<code>syslog(3C)</code> , <code>termios(3C)</code> , <code>ldterm(7M)</code> , <code>termio(7I)</code> , <code>msglog(7D)</code> , <code>wscons(7D)</code>
NOTES	<p>In contrast to pre-SunOS 5.0 releases, it is no longer possible to redirect I/O intended for <code>/dev/console</code> to some other device. Instead, redirection now applies to the workstation console device using a revised programming interface (see <code>wscons(7D)</code>). Since the system console is normally configured to be the work station console, the overall effect is largely unchanged from previous releases.</p> <p>See <code>wscons(7D)</code> for detailed descriptions of control sequence syntax, ANSI control functions, control character functions and escape sequence functions.</p>

NAME	cpr – Suspend and resume module
SYNOPSIS	<code>/platform/'uname -m'/kernel/misc/cpr</code>
DESCRIPTION	<p>The <code>cpr</code> module is a loadable module used to suspend and resume the entire system. You may wish to suspend a system to save power or to power off temporarily for transport. The <code>cpr</code> module should not be used in place of a normal shutdown when performing any hardware reconfiguration or replacement. In order for the resume operation to succeed, it is important that the hardware configuration remain the same. When the system is suspended, the entire system state is preserved in non-volatile storage until a resume operation is conducted.</p> <p><code>dtpower(1M)</code> or <code>power.conf(4)</code> are used to configure the suspend-resume feature.</p> <p>The speed of suspend and resume operations can range from 15 seconds to several minutes, depending on the system speed, memory size, and load.</p> <p>During resume operation, the <code>SIGTHAW</code> signal is sent to all processes to allow them to do any special processing in response to suspend-resume operation. Normally applications are not required to do any special processing because of suspend-resume, but some specialized processes can use <code>SIGTHAW</code> to restore the state prior to suspend. For example, X can refresh the screen in response to <code>SIGTHAW</code>.</p> <p>In some cases the <code>cpr</code> module may be unable to perform the suspend operation. If a system contains additional devices outside the standard shipped configuration, it is possible that device drivers for these additional devices might not support suspend-resume operations. In this case, the suspend fails and an error message is displayed. These devices must be removed or their device drivers unloaded for the suspend operation to succeed. Contact the device manufacturer to obtain a new version of device driver that supports suspend-resume.</p> <p>A suspend may also fail when devices or processes are performing critical or time-sensitive operations (such as realtime operations). The system will remain in its current running state. Messages reporting the failure will be displayed on the console and status returned to the caller. Once the system is successfully suspended the resume operation will succeed, barring external influences such as a hardware reconfiguration.</p> <p>Some network-based applications may fail across a suspend and resume cycle. This largely depends on the underlying network protocol and the applications involved. In general, applications that retry and automatically reestablish connections will continue to operate transparently on a resume operation; those applications that do not will likely fail.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

cpr(7)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcpr
Interface stability	Unstable

SEE ALSO `dtpower(1M)` (OpenWindows Reference Manual), `pmconfig(1M)`, `uadmin(1M)`, `uadmin(2)`, `power.conf(4)`, `attributes(5)`

Using Power Management

Writing Device Drivers

NOTES Certain device operations such as tape and floppy disk activities are not resumable due to the nature of removable media. These activities are detected at suspend time, and must be stopped before the suspend operation will complete successfully.

Suspend-resume is currently supported only on a limited set of hardware platforms. Please see the book *Using Power Management* for a complete list of platforms that support system Power Management. See `uname(2)` to programatically determine if the machine supports suspend-resume.

NAME	cpuid – CPU identification driver
SYNOPSIS	<code>/dev/cpu/self/cpuid</code>
DESCRIPTION	
SPARC and x86 system	This device provides implementation-private information via ioctls about various aspects of the implementation to Solaris libraries and utilities.
x86 systems only	<p>This device also provides a file-like view of the namespace and return values of the x86 cpuid instruction. The cpuid instruction takes a single 32-bit integer function code, and returns four 32-bit integer values corresponding to the input value that describe various aspects of the capabilities and configuration of the processor.</p> <p>The API for the character device consists of using the seek offset to set the function code value, and using a <code>read(2)</code> or <code>pread(2)</code> of 16 bytes to fetch the four 32-bit return values of the instruction in the order <code>%eax</code>, <code>%ebx</code>, <code>%ecx</code> and <code>%edx</code>.</p> <p>No data can be written to the device. Like the cpuid instruction, no special privileges are required to use the device.</p> <p>The device is useful to enable low-level configuration information to be extracted from the CPU without having to write any assembler code to invoke the cpuid instruction directly. It also allows the kernel to attempt to correct any erroneous data returned by the instruction (prompted by occasional errors in the information exported by various processor implementations over the years).</p> <p>See the processor manufacturers documentation for further information about the syntax and semantics of the wide variety of information available from this instruction.</p>
EXAMPLE	<p>This example allows you to determine if the current x86 processor supports "long mode," which is a necessary (but not sufficient) condition for running the 64-bit Solaris kernel on the processor.</p> <pre> /* #include <sys/types.h> #include <sys/stat.h> #include <fcntl.h> #include <unistd.h> #include <string.h> #include <errno.h> #include <stdio.h> static const char devname[] = "/dev/cpu/self/cpuid"; /*ARGSUSED*/ int main(int argc, char *argv[]) { struct { uint32_t r_eax, r_ebx, r_ecx, r_edx; } _r, *rp = &_amp_r; </pre>

cpuid(7D)

```

int d;
char *s;

if ((d = open(devname, O_RDONLY)) == -1) {
    perror(devname);
    return (1);
}

if (pread(d, rp, sizeof (*rp), 0) != sizeof (*rp)) {
    perror(devname);
    goto fail;
}

s = (char *)&rp->r_ebx;
if (strncmp(s, "Auth" "cAMD" "enti", 12) != 0 &&
    strncmp(s, "Genu" "ntel" "ineI", 12) != 0)
    goto fail;

if (pread(d, rp, sizeof (*rp), 0x80000001) == sizeof (*rp)) {
    /*
     * Read extended feature word; check bit 29
     */
    (void) close(d);
    if ((rp->r_edx >> 29) & 1) {
        (void) printf("processor supports long mode\n");
        return (0);
    }
}

fail:
    (void) close(d);
    return (1);
}

```

- ERRORS**
- ENXIO** Results from attempting to read data from the device on a system that does not support the CPU identification interfaces
- EINVAL** Results from reading from an offset larger than `UINT_MAX`, or attempting to read with a size that is not multiple of 16 bytes.

FILES `/dev/cpu/self/cpuid`
 Provides access to CPU identification data.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWckr
Interface Stability	Evolving

SEE ALSO `psrinfo(1M)`, `prtconf(1M)`, `pread(2)`, `read(2)`, `attributes(5)`

- NAME** | ctfs – contract file system
- DESCRIPTION** | The `ctfs` filesystem is the interface to the contract sub-system. `ctfs` is mounted during boot at `/system/contract`. For information on contracts and the contents of this filesystem, see `contract(4)`.
- FILES** | `/system/contract` | Mount point for the `ctfs` file system
- ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWckr

SEE ALSO | `contract(4)`, `vfstab(4)`, `attributes(5)`, `smf(5)`

ctsmc(7D)

NAME | ctsmc – System Management Controller driver

DESCRIPTION | The `ctsmc` system management controller driver is a multithreaded, loadable, clonable STREAMS hardware driver that supports communication with the system management controller device on SUNW,NetraCT-410, SUNW,NetraCT-810 and SUNW,Netra-CP2300 platforms.

The `smc` device provides a Keyboard Controller Style (KCS) interface as described in the *Intelligent Platform Management Interface (IPMI) Version 1.5* specification. The `ctsmc` driver enables user-land and kernel-land clients to access services provided by `smc` hardware.

FILES | `/dev/ctsmc`
| `ctsmc` special character device

| `/platform/sun4u/kernel/drv/sparcv9/ctsmc`
| 64 bit ELF kernel driver

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Availability	SUNWcar.u

SEE ALSO | `attributes(5)`

| *STREAMS Programmers Guide*

| *Writing Device Drivers*

| *Intelligent Platform Management Interface (IPMI). Version 1.5 – PICMIG, February, 2001*

NAME	cvc – virtual console driver
DESCRIPTION	<p>The <code>cvc</code> virtual console driver is a STREAMS-based pseudo driver that supports the network console. The <code>cvc</code> driver interfaces with <code>console(7D)</code>.</p> <p>Logically, the <code>cvc</code> driver sits below the <code>console</code> driver. It redirects console output to the <code>cvcredir(7D)</code> driver if a network console connection is active. If a network console connection is not active, it redirects console output to an internal hardware interface.</p> <p>The <code>cvc</code> driver receives console input from <code>cvcredir</code> and internal hardware and passes it to the process associated with <code>/dev/console</code>.</p>
NOTES	The <code>cvc</code> facility supersedes the SunOS <code>wscons(7D)</code> facility, which should <i>not</i> be used in conjunction with <code>cvc</code> . The <code>wscons</code> driver is useful for systems with directly attached consoles (frame buffers and keyboards), but is not useful with platforms using <code>cvc</code> , which have no local keyboard or frame buffer.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Sun Enterprise 10000 servers, Sun Fire 15000 servers
Availability	SUNWcvc.u

SEE ALSO	<p><code>cvcd(1M)</code>, <code>attributes(5)</code>, <code>console(7D)</code>, <code>cvcredir(7D)</code>, <code>wscons(7D)</code></p> <p><i>Sun Enterprise 10000 SSP Reference Manual</i></p> <p><i>Sun System Management Services (SMS) Reference Manual</i></p>
-----------------	--

cvcredir(7D)

NAME cvcredir – virtual console redirection driver

DESCRIPTION The `cvcredir` virtual console redirection driver is a STREAMS-based pseudo driver that supports the network console provided on some platforms. The `cvcredir` driver interfaces with the virtual console driver `cvc(7D)`, and the virtual console daemon, `cvcd(1M)`.

The `cvcredir` driver receives console output from `cvc` and passes it to `cvcd`. It receives console input from `cvcd` and passes it to `cvc`.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Sun Enterprise 10000 servers, Sun Fire 15K servers
Availability	SUNWcvc.u

SEE ALSO `cvcd(1M)`, `attributes(5)`, `console(7D)`, `cvc(7D)`

Sun Enterprise 10000 SSP Reference Manual

Sun System Management Services (SMS) Reference Manual

NAME	dad – driver for IDE disk devices	
SYNOPSIS	dad@ <i>target</i> , <i>lun</i> : <i>partition</i>	
DESCRIPTION	<p>This driver handles the ide disk drives on SPARC platforms. The type of disk drive is determined using the ATA IDE identify device command and by reading the volume label stored on block 0 of the drive. The volume label describes the disk geometry and partitioning; it must be present or the disk cannot be mounted by the system.</p> <p>The block-files access the disk using the system's normal buffering mechanism and are read and written without regard to physical disk records. There is also a "raw" interface that provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in one I/O operation; raw I/O is therefore considerably more efficient when many bytes are transmitted. The names of the block files are found in <code>/dev/dsk</code>; the names of the raw files are found in <code>/dev/rdsk</code>.</p> <p>I/O requests to the raw device must be aligned on a 512-byte (<code>DEV_BSIZE</code>) boundary and must have a length that is a multiple of 512 bytes. Requests which do not meet the restrictions will cause the driver to return an <code>EINVAL</code> error. I/O requests to the block device have no alignment or length restrictions.</p>	
Device Statistics Support	<p>Each device maintains I/O statistics both for the device and for each partition allocated on that device. For each device/partition, the driver accumulates reads, writes, bytes read, and bytes written. The driver also takes hi-resolution time stamps at queue entry and exit points, which facilitates monitoring the residence time and cumulative residence-length product for each queue.</p> <p>Each device also has error statistics associated with it. These must include counters for hard errors, soft errors and transport errors. Other data may be implemented as required.</p>	
FILES	<code>/dev/dsk/cntndnsn</code>	block files
	<code>/dev/rdsk/cntndnsn</code>	raw files
	where:	
	<code>cn</code>	controller <i>n</i>
	<code>tn</code>	IDE target id <i>n</i> (0-3)
	<code>dn</code>	Always 0.
	<code>sn</code>	partition <i>n</i> (0-7)
	The target ide numbers are assigned as:	
	0	Master disk on Primary channel.
	1	Slave disk on Primary channel.
	2	Master disk on Secondary channel
	3	Slave disk on Secondary channel.

dad(7D)

IOCTLS	Refer to dkio(7I) .																
ERRORS	<table><tr><td>EACCES</td><td>Permission denied.</td></tr><tr><td>EBUSY</td><td>The partition was opened exclusively by another thread.</td></tr><tr><td>EFAULT</td><td>The argument was a bad address.</td></tr><tr><td>EINVAL</td><td>Invalid argument.</td></tr><tr><td>EIO</td><td>An I/O error occurred.</td></tr><tr><td>ENOTTY</td><td>This indicates that the device does not support the requested ioctl function.</td></tr><tr><td>ENXIO</td><td>During opening, the device did not exist.</td></tr><tr><td>EROFS</td><td>The device is a read-only device.</td></tr></table>	EACCES	Permission denied.	EBUSY	The partition was opened exclusively by another thread.	EFAULT	The argument was a bad address.	EINVAL	Invalid argument.	EIO	An I/O error occurred.	ENOTTY	This indicates that the device does not support the requested ioctl function.	ENXIO	During opening, the device did not exist.	EROFS	The device is a read-only device.
EACCES	Permission denied.																
EBUSY	The partition was opened exclusively by another thread.																
EFAULT	The argument was a bad address.																
EINVAL	Invalid argument.																
EIO	An I/O error occurred.																
ENOTTY	This indicates that the device does not support the requested ioctl function.																
ENXIO	During opening, the device did not exist.																
EROFS	The device is a read-only device.																
SEE ALSO	format(1M) , mount(1M) , lseek(2) , read(2) , write(2) , driver.conf(4) , vfstab(4) , dkio(7I) X3T10 ATA-4 specifications.																
DIAGNOSTICS	<p>Command: <i><number></i>, Error:<i><number></i>, Status:<i><number></i> Indicates that the command failed with an error and provides status register contents. Where <i><number></i> is a hexadecimal value.</p> <p>offline The driver has decided that the target disk is no longer there.</p> <p>disk ok The target disk is now responding again.</p> <p>corrupt label - bad geometry The disk label is corrupted.</p> <p>corrupt label - label checksum failed The disk label is corrupted.</p> <p>corrupt label - wrong magic number The disk label is corrupted.</p> <p>disk not responding to selection The target disk is not responding.</p> <p>i/o to invalid geometry The geometry of the drive could not be established.</p> <p>incomplete read/write - retrying/giving up There was a residue after the command completed normally.</p> <p>no bp for disk label A bp with consistent memory could not be allocated.</p> <p>no memory for disk label Free memory pool exhausted.</p>																

dad(7D)

```
ATA transport failed: reason 'nnnn': {retrying|giving}
  The host adapter has failed to transport a command to the target for the reason
  stated. The driver will either retry the command or, ultimately, give up.

corrupt label - wrong magic number
  The disk label is corrupted.

corrupt label - label checksum failed
  The disk label is corrupted.

corrupt label - bad geometry
  The disk label is corrupted.

no mem for property
  Free memory pool exhausted.

transport rejected (<n>)
  Host adapter driver was unable to accept a command.

Device Fault
  There has been a Device Fault - reason for such error is vendor specific.
```

daplt(7D)

NAME	daplt – Tavor uDAPL service driver						
SYNOPSIS	daplt@0:daplt						
DESCRIPTION	<p>The <code>daplt</code> module is the driver component of the uDAPL service provider for Tavor which implements the provider functions of the <i>uDAPL Specification 1.2</i> described under <code>libdat(3LIB)</code>.</p> <p>The <code>daplt</code> module is a child of the IB nexus driver, <code>ib(7D)</code>, and layers on top of the Solaris kernel IB Transport Layer, <code>ibt1(7D)</code>. The <code>daplt</code> driver uses the InfiniBand Transport Framework (IBTF). (See <code>ibt1(7D)</code>, <code>ibcm(7D)</code>, and <code>ib(7D)</code> to access privileged IB VERBS.)</p> <p>The <code>daplt</code> driver copies out various HCA H/W object reference handles, including working and completion queues and User Access Region registers, to its own uDAPL service provider library for Tavor. The library can refer back to these object handles and use them to <code>mmap(2)</code> in the mapping of these H/W queues and registers from the Tavor HCA driver, <code>tavor(7D)</code>. This process enables time-critical non-privileged IB VERBS such as send/receive work elements, RDMA read/write and memory window bind, to be invoked in the userland library and performed directly by the firmware or hardware. As a result, OS and network stack are bypassed, achieving true zero data copy with the lowest possible latency.</p>						
FILES	<p><code>/kernel/drv/sparcv9/daplt</code> 64-bit SPARC ELF kernel driver</p> <p><code>/kernel/drv/daplt</code> 32-bit x86 ELF kernel driver</p> <p><code>/kernel/drv/amd64/daplt</code> 64-bit x86 ELF kernel driver</p> <p><code>/kernel/drv/daplt.conf</code> driver configuration file</p> <p><code>/dev/daplt</code> special character device.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>SPARC, x86</td></tr><tr><td>Availability</td><td>SUNWdapltr</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC, x86	Availability	SUNWdapltr
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	SPARC, x86						
Availability	SUNWdapltr						
SEE ALSO	<code>mmap(2)</code> , <code>libdat(3LIB)</code> , <code>driver.conf(4)</code> , <code>attributes(5)</code> , <code>ib(7D)</code> , <code>ibcm(7D)</code> , <code>ibdm(7D)</code> , <code>ibt1(7D)</code> , <code>tavor(7D)</code> <i>uDAPL Specification 1.2</i>						

NAME	dbri – Dual Basic Rate ISDN and audio Interface
DESCRIPTION	<p>The dbri device uses the T5900FC Dual Basic Rate ISDN Interface (DBRI) and Multimedia Codec chips to implement the audio device interface. This interface is described fully in the <code>audio(7I)</code> manual page.</p> <p>Applications that open <code>/dev/audio</code> may use the <code>AUDIO_GETDEV</code> ioctl to determine which audio device is being used. The dbri driver will return the string "SUNW, dbri" in the <code>name</code> field of the <code>audio_device</code> structure. The <code>Version</code> field will contain "e" and the <code>Config</code> field will contain one of the following values: "isdn_b" on an ISDN B channel stream, "speakerbox" on a <code>/dev/audio</code> stream associated with a SpeakerBox, and lastly "onboard1" on a <code>/dev/audio</code> stream associated with the onboard Multimedia Codec.</p> <p>The <code>AUDIO_SETINFO</code> ioctl controls device configuration parameters. When an application modifies the <code>record.buffer_size</code> field using the <code>AUDIO_SETINFO</code> ioctl, the driver will constrain it to be non-zero and a multiple of 16 bytes, up to a maximum of 8176 bytes.</p>
Audio Interfaces	<p>The SpeakerBox audio peripheral is available for connection to the SpeakerBox Interface (SBI) port of most dbri equipped systems and provides an integral monaural speaker as well as stereo line out, stereo line in, stereo headphone, and monaural microphone connections. The headset output level is adequate to power most headphones, but may be too low for some external speakers. Powered speakers or an external amplifier may be used with both the headphone and line out ports.</p> <p>The Sun Microphone is recommended for normal desktop audio recording. When the Sun Microphone is used in conjunction with the SpeakerBox, the microphone battery is bypassed. Other audio sources may be recorded by connecting their line output to the SpeakerBox line input (audio sources may also be connected from their headphone output if the volume is adjusted properly).</p>
ISDN Interfaces	<p>The DBRI controller offers two Basic Rate ISDN (BRI) interfaces. One is a BRI Terminal Equipment (TE) interface and the other is a BRI Network Termination (NT) interface.</p> <p>The NT connector is switched by a relay so that when system power is not available or when software is not accessing the NT port, the TE and NT connectors are electrically connected and devices plugged into the NT port will be on the same BRI passive bus.</p>
Audio Data Formats for the Multimedia Codec/SpeakerBox	<p>The dbri device supports the audio formats listed in the following table. When the device is open for simultaneous play and record, the input and output data formats must match.</p>

Supported Audio Data Formats

Sampe Rate	Encoding	Precision	Channels
8000 Hz	mu-law or A-law	8	1

dbri(7D)

Supported Audio Data Formats			
Sampe Rate	Encoding	Precision	Channels
9600 Hz	mu-law or A-law	8	1
11025 Hz	mu-law or A-law	8	1
16000 Hz	mu-law or A-law	8	1
18900 Hz	mu-law or A-law	8	1
22050 Hz	mu-law or A-law	8	1
32000 Hz	mu-law or A-law	8	1
37800 Hz	mu-law or A-law	8	1
44100 Hz	mu-law or A-law	8	1
48000 Hz	mu-law or A-law	8	1
8000 Hz	linear	16	1 or 2
9600 Hz	linear	16	1 or 2
11025 Hz	linear	16	1 or 2
16000 Hz	linear	16	1 or 2
18900 Hz	linear	16	1 or 2
22050 Hz	linear	16	1 or 2
32000 Hz	linear	16	1 or 2
37800 Hz	linear	16	1 or 2
44100 Hz	linear	16	1 or 2
48000 Hz	linear	16	1 or 2

**Audio Data
Formats for BRI
Interfeces**

ISDN channels implement a subset of audio semantics. The preferred ioctls for querying or setting the format of a BRI channel are `ISDN_GET_FORMAT`, `ISDN_SET_FORMAT`, and `ISDN_SET_CHANNEL`. In particular, there is no audio format described in `audio(7I)` that covers HDLC or transparent data. The `dbri` driver maps HDLC and transparent data to `AUDIO_ENCODING_NONE`. ISDN D-channels are always configured for HDLC encoding of data. The programmer should interpret an *encoding* value of `AUDIO_ENCODING_NONE` as an indication that the *fd* is not being used to transfer audio data.

B-channels can be configured for mu-law (as in the Greek letter mu), A-law, or HDLC encoding of data. The mu-law and A-law formats are always at 8000 Hz, 8-bit, mono. Although a BRI H-channel is actually 16 bits wide at the physical layer and the 16-bit sample occurs at 8 kHz, the HDLC encoding always presents the data in 8-bit

dbri(7D)

quantities. Therefore, 56 bit-per-second (bps), 64 bps, and 128 bps formats are all presented to the programmer as 8-bit wide, mono, AUDIO_ENCODING_NONE format streams at different sample rates. A line rate of 56kbps results in a 8-bit sample rate of 7000 Hz. If the bit stuffing and un-stuffing of HDLC were taken into account, the data rate would be slightly less.

For the sake of compatibility, AUDIO_GETINFO will return one of the following on a ISDN channel:

BRI Audio Data Formats			
Sample Rate	Encoding	Precision	Channels
8000 Hz	mu-law or A-law	8	1
-	AUDIO_ENCODING_NONE	-	-

ISDN_GET_FORMAT will return one of the following for an ISDN channel:

BRI Audio Data Formats					
Mode	Sample Rate	Encoding	Precision	# Ch	Available on
HDLC	2000 Hz	NONE	8	1	D
HDLC	7000 Hz	NONE	8	1	B1,B2
HDLC	8000 Hz	NONE	8	1	B1,B2
HDLC	16000 Hz	NONE	8	1	B1,B2
TRANS	8000 Hz	mu-law	8	1	B1,B2
TRANS	8000 Hz	A-law	8	1	B1,B2
TRANS	8000 Hz	NONE	8	1	B1,B2
TRANS	8000 Hz	NONE	16	1	B1 only

In the previous table:

HDLC = ISDN_MODE_HDLC TRANS = ISDN_MODE_TRANSPARENT

Audio Ports

Audio ports are not relevant to ISDN D or B channels.

The *record.avail_ports* and *play.avail_ports* fields of the *audio_info* structure report the available input and output ports. The *dbri* device supports two input ports, selected by setting the *record.port* field to either AUDIO_MICROPHONE or AUDIO_LINE_IN. The *play.port* field may be set to any combination of AUDIO_SPEAKER, AUDIO_HEADPHONE, and AUDIO_LINE_OUT by OR'ing the desired port names together.

dbri(7D)

Sample Granularity	Since the <code>dbri</code> device manipulates buffers of audio data, at any given time the reported input and output sample counts will vary from the actual sample count by no more than the size of the buffers it is transferring. Programs should, in general, not rely on absolute accuracy of the <code>play.samples</code> and <code>record.samples</code> fields of the <code>audio_info</code> structure.
Audio Status Change Notification	As described in <code>audio(7I)</code> , it is possible to request asynchronous notification of changes in the state of an audio device. The DBRI driver extends this to the ISDN B channels by sending the signal up the data channel instead of the control channel. Asynchronous notification of events on a B-channel only occurs when the channel is in a transparent data mode. When the channel is in HDLC mode, no such notification will take place.
ERRORS	In addition to the errors described in <code>audio(7I)</code> , an <code>open()</code> will fail if: ENODEV The driver is unable to communicate with the SpeakerBox, possibly because it is currently not plugged in.
FILES	The physical device names are very system dependent and are rarely used by programmers. For example: <code>/devices/sbus@1,f8000000/SUNW,DBRIe@1,10000:te,b2.</code> The programmer should instead use the generic device names listed below: <code>/dev/audio</code> symlink to the system's primary audio device, not necessarily a <code>dbri</code> based audio device <code>/dev/audioctrl</code> control device for the above audio device <code>/dev/sound/0*</code> represents the first audio device on the system and is not necessarily based on <code>dbri</code> or SpeakerBox <code>/dev/sound/0</code> first audio device in the system <code>/dev/sound/0ctrl</code> audio control for above device <code>/dev/isdn/0/*</code> represents the first ISDN device on the system and any associated interfaces. This device is not necessarily based on <code>dbri</code> . <code>/dev/isdn/0/te/mgt</code> TE management device <code>/dev/isdn/0/te/d</code> TE D channel <code>/dev/isdn/0/te/b1</code> TE B1 channel <code>/dev/isdn/0/te/b2</code> TE B2 channel <code>/dev/isdn/0/nt/mgt</code> NT management device <code>/dev/isdn/0/nt/d</code> NT D channel <code>/dev/isdn/0/nt/b1</code> NT B1 channel <code>/dev/isdn/0/nt/b2</code> NT B2 channel

dbri(7D)

/dev/isdn/0/aux/0 SpeakerBox or onboard Multimedia Codec
/dev/isdn/0/aux/0ctl Control device for SpeakerBox or onboard Multimedia
Codec
/usr/share/audio Audio files

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO `ioctl(2)`, `attributes(5)`, `audio(7I)`, `isdnic(7I)`, `streamio(7I)`

AT&T Microelectronics data sheet for the T5900FC Sun Dual Basic Rate ISDN Interface.

Crystal Semiconductor, Inc., data sheet for the CS4215 16-Bit, 48 kHz, Multimedia Audio Codec Publication number DS76PP5.

NOTES Due to hardware restrictions, it is impossible to reduce the record gain to 0. A valid input signal is still received at the lowest gain setting the Multimedia Codec allows. For security reasons, the `dbri` driver disallows a record gain value of 0. This is to provide feedback to the user that such a setting is not possible and that a valid input signal is still being received. An attempt to set the record gain to 0 will result in the lowest possible non-zero gain. The `audio_info` structure will be updated with this value when the `AUDIO_SETINFO` `ioctl` returns.

BUGS When a DBRI channel associated with the SpeakerBox Interface underruns, DBRI may not always repeat the last sample but instead could repeat more than one sample. This behavior can result in a tone being generated by an audio device connected to the SBI port.

Monitor STREAMs connected to a B1 channel on either the TE or NT interface do not work because of a DBRI hardware problem. The device driver disallows the creation of such monitors.

devfs(7FS)

NAME	devfs – Devices file system
DESCRIPTION	<p>The <code>devfs</code> filesystem manages a name space of all devices under the Solaris operating environment and is mounted during boot on the <code>/devices</code> name space.</p> <p>The <code>/devices</code> name space is dynamic and reflects the current state of accessible devices under the Solaris operating environment. The names of all attached device instances are present under <code>/devices</code>.</p> <p>The content under <code>/devices</code> is under the exclusive control of the <code>devfs</code> filesystem and cannot be changed.</p> <p>The system may be configured to include a device in one of two ways:</p> <p>By means of dynamic reconfiguration (DR), using, for example, <code>cfgadm(1M)</code>.</p> <p>For devices driven by <code>driver.conf(4)</code> enumeration, edit the <code>driver.conf</code> file to add a new entry, then use <code>update_drv(1M)</code> to cause the system to re-read the <code>driver.conf</code> file and thereby enumerate the instance.</p> <p>The device may be attached through a number of system calls and programs, including <code>open(2)</code>, <code>stat(2)</code> and <code>ls(1)</code>. During device attach, the device driver typically creates minor nodes corresponding to the device via <code>ddi_create_minor_node(9F)</code>. If the attach is successful, one or more minor nodes referring to the device are created under <code>/devices</code>.</p> <p>Operations like <code>mknod(2)</code>, <code>mkdir(2)</code> and <code>creat(2)</code> are not supported in <code>/devices</code>.</p>
FILES	<code>/devices</code> Mount point for <code>devfs</code> file system
SEE ALSO	<code>devfsadm(1M)</code> , <code>vfstab(4)</code> , <code>attach(9E)</code>
NOTES	<p>The <code>/devices</code> name space cannot be unmounted.</p> <p>All content at or below the <code>/devices</code> name space is an implementation artifact and subject to incompatible change or removal without notification.</p>

- NAME** devinfo – device information driver
- DESCRIPTION** The devinfo driver is a private mechanism used by the libdevinfo(3LIB) interfaces to access kernel device configuration data and to guarantee data consistency.
- FILES** /devices/pseudo/devinfo@0:devinfo
- ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Stability Level	Private

SEE ALSO libdevinfo(3LIB), attributes(5)

Writing Device Drivers

dkio(7I)

NAME	dkio – disk control operations
SYNOPSIS	<pre>#include <sys/dkio.h> #include <sys/vtoc.h></pre>
DESCRIPTION	Disk drivers support a set of <code>ioctl(2)</code> requests for disk controller, geometry, and partition information. Basic to these <code>ioctl()</code> requests are the definitions in <code><sys/dkio.h></code> .
IOCTLS	<p>The following <code>ioctl()</code> requests set and/or retrieve the current disk controller, partitions, or geometry information on all architectures:</p> <p>DKIOCIINFO</p> <p>The argument is a pointer to a <code>dk_cinfo</code> structure (described below). This structure tells the controller-type and attributes regarding bad-block processing done on the controller.</p> <pre>/* * Structures and definitions for disk I/O control commands */ #define DK_DEVLEN 16 /* device name max length, */ /* including unit # and NULL */ /* Used for controller info */ struct dk_cinfo { char dki_cname[DK_DEVLEN]; /* controller name */ /*(no unit #)*/ ushort_t dki_ctype; /* controller type */ ushort_t dki_flags; /* flags */ ushort_t dki_cnum; /* controller number */ uint_t dki_addr; /* controller address */ uint_t dki_space; /* controller bus type */ uint_t dki_prio; /* interrupt priority */ uint_t dki_vec; /* interrupt vector */ char dki_dname[DK_DEVLEN]; /* drive name (no unit #) */ uint_t dki_unit; /* unit number */ uint_t dki_slave; /* slave number */ ushort_t dki_partition; /* partition number */ ushort_t dki_maxtransfer; /* maximum transfer size */ /* in DEV_BSIZE */ }; /* * Controller types */ #define DKC_UNKNOWN 0 #define DKC_CDROM 1 /* CD-ROM, SCSI or other */ #define DKC_WDC2880 2 #define DKC_XXX_0 3 /* unassigned */ #define DKC_XXX_1 4 /* unassigned */ #define DKC_DSD5215 5 #define DKC_ACB4000 7 #define DKC_XXX_2 9 /* unassigned */ #define DKC_NCRFLOPPY 10 #define DKC_SMSFLOPPY 12 #define DKC_SCSI_CCS 13 /* SCSI CCS compatible */</pre>

```

#define DKC_INTEL82072  14      /* native floppy chip */
#define DKC_MD          16      /* meta-disk (virtual-disk) */
                                /* driver */
#define DKC_INTEL82077  19      /* 82077 floppy disk */
                                /* controller */
#define DKC_DIRECT      20      /* Intel direct attached */
                                /* device (IDE) */
#define DKC_PCPCIA_MEM  21      /* PCMCIA memory disk-like */
                                /* type */
#define DKC_PCPCIA_ATA  22      /* PCMCIA AT Attached type */

/*
 * Sun reserves up through 1023
 */

#define DKC_CUSTOMER_BASE 1024

/*
 * Flags
 */

#define DKI_BAD144      0x01     /* use DEC std 144 */
                                /* bad sector fwding */
#define DKI_MAPTRK     0x02     /* controller does */
                                /* track mapping */
#define DKI_FMTTRK     0x04     /* formats only full
                                /* track at a time*/
#define DKI_FMTVOL     0x08     /* formats only full */
                                /* volume at a time*/
#define DKI_FMTCYL     0x10     /* formats only full */
                                /* cylinders at a time*/
#define DKI_HEXUNIT    0x20     /* unit number printed as */
                                /* 3 hexdigits */
#define DKI_PCPCIA_PFD 0x40     /* PCMCIA pseudo-floppy */
                                /* memory card */

```

DKIOCGAPART

The argument is a pointer to a `dk_allmap` structure (described below). This `ioctl()` gets the controller's notion of the current partition table for disk drive.

DKIOCSAPART

The argument is a pointer to a `dk_allmap` structure (described below). This `ioctl()` sets the controller's notion of the partition table without changing the disk itself.

```

/*
 * Partition map (part of dk_label)
 */ struct dk_map {
    daddr_t dkl_cylno;    /* starting cylinder */
    daddr_t dkl_nblk;    /* number of blocks */
};

/*
 * Used for all partitions
 */
struct dk_map {
struct dk_allmap {

```

dkio(7I)

```
struct dk_map dka_map[NDKMAP];
};
```

DKIOCGGGEOM

The argument is a pointer to a `dk_geom` structure (described below). This `ioctl()` gets the controller's notion of the current geometry of the disk drive.

DKIOCSGGEOM

The argument is a pointer to a `dk_geom` structure (described below). This `ioctl()` sets the controller's notion of the geometry without changing the disk itself.

DKIOCGVTOC

The argument is a pointer to a `vtoc` structure (described below). This `ioctl()` returns the device's current volume table of contents (VTOC.)

DKIOCSVTOC

The argument is a pointer to a `vtoc` structure (described below). This `ioctl()` changes the VTOC associated with the device.

```
struct partition {
ushort_t    p_tag;           /* ID tag of partition */
ushort_t    p_flag;         /* permission flags */
daddr_t     p_start;        /* start sector of partition */
long        p_size;         /* # of blocks in partition */
};
```

If `DKIOCSVTOC` is used with a floppy diskette, the `p_start` field must be the first sector of a cylinder. To compute the number of sectors per cylinder, multiply the number of heads by the number of sectors per track.

```
struct vtoc {
unsigned long    v_bootinfo[3];           /* info needed by mboot
                                           /* (unsupported)*/
unsigned long    v_sanity;                /* to verify vtoc sanity */
unsigned long    v_version;               /* layout version */
char             v_volume[LEN_DKL_VVOL];  /* volume name */
ushort_t         v_sectorsz;              /*
    sector size in bytes*/
ushort_t         v_nparts;                /*
    number of partitions*/
unsigned long    v_reserved[10];          /* free space */
struct partition v_part[V_NUMPAR];        /* partition headers*/
time_t          timestamp[V_NUMPAR];      /* partition timestamp
                                           /* (unsupported)*/
char            v_asciilabel[LEN_DKL_ASCII]; /* compatibility */
};

/*
 * Partition permission flags
 */

#define V_UNMNT    0x01 /* Unmountable partition */
#define V_RDONLY  0x10 /* Read only */
```



```

/*
 * Partition identification tags
 */

#define V_UNASSIGNED 0x00 /* unassigned partition */
#define V_BOOT 0x01 /* Boot partition */
#define V_ROOT 0x02 /* Root filesystem */
#define V_SWAP 0x03 /* Swap filesystem */
#define V_USR 0x04 /* Usr filesystem */
#define V_BACKUP 0x05 /* full disk */
#define V_VAR 0x07 /* Var partition */
#define V_HOME 0x08 /* Home partition */
#define V_ALTSCTR 0x09 /* Alternate sector partition */

```

DKIOCEJECT

If the drive supports removable media, this `ioctl()` requests the disk drive to eject its disk.

DKIOCREMOVABLE

The argument to this `ioctl()` is an integer. After successful completion, this `ioctl()` sets that integer to a non-zero value if the drive in question has removable media. If the media is not removable, that integer is set to 0.

DKIOCSTATE

This `ioctl()` blocks until the state of the drive, inserted or ejected, is changed. The argument is a pointer to a `dkio_state`, enum, whose possible enumerations are listed below. The initial value should be either the last reported state of the drive, or `DKIO_NONE`. Upon return, the enum pointed to by the argument is updated with the current state of the drive.

```

enum dkio_state {
DKIO_NONE,          /* Return disk's current state */
DKIO_EJECTED,      /* Disk state is 'ejected' */
DKIO_INSERTED      /* Disk state is 'inserted' */
};

```

DKIOCLLOCK

For devices with removable media, this `ioctl()` requests the disk drive to lock the door.

DKIOCUNLOCK

For devices with removable media, this `ioctl()` requests the disk drive to unlock the door.

DKIOCGMEDIAINFO

The argument to this `ioctl()` is a pointer to a `dk_minfo` structure. The structure indicates the type of media or the command set profile used by the drive to operate on the media. The `dk_minfo` structure also indicates the logical media blocksize the drive uses as the basic unit blocksize of operation and the raw formatted capacity of the media in number of logical blocks.

```

/*
 * Used for media info or profile info

```

dkio(7I)

```
*/
struct dk_minfo {
    uint_t      dki_media_type; /* Media type or profile info */
    uint_t      dki_lbsize;     /* Logical blocksize of media */
    diskaddr_t  dki_capacity;   /* Capacity as # of dki_lbsize blks */
};
/*
 * Media types or profiles known
 */
#define DK_UNKNOWN          0x00 /* Media inserted - type unknown */

/*
 * SFF 8090 Specification Version 3, media types 0x01 - 0xfffe are retained to
 * maintain compatibility with SFF8090. The following define the
 * optical media type.
 */
#define DK_MO_ERASABLE      0x03 /* MO Erasable */
#define DK_MO_WRITEONCE    0x04 /* MO Write once */
#define DK_AS_MO           0x05 /* AS MO */
#define DK_CDROM           0x08 /* CDROM */
#define DK_CDR             0x09 /* CD-R */
#define DK_CDRW            0x0A /* CD-RW */
#define DK_DVDROM          0x10 /* DVD-ROM */
#define DK_DVDR            0x11 /* DVD-R */
#define DK_DVDRAM          0x12 /* DVD_RAM or DVD-RW */

/*
 * Media types for other rewritable magnetic media
 */
#define DK_FIXED_DISK      0x10001 /* Fixed disk SCSI or otherwise */
#define DK_FLOPPY          0x10002 /* Floppy media */
#define DK_ZIP             0x10003 /* IOMEGA ZIP media */
#define DK_JAZ             0x10004 /* IOMEGA JAZ media */
```

If the media exists and the host can obtain a current profile list, the command succeeds and returns the `dk_minfo` structure with data representing that media.

If there is no media in the drive, the command fails and the host returns an `ENXIO` error, indicating that it cannot gather the information requested.

If the profile list is not available, the host attempts to identify the media-type based on the available information.

If identification is not possible, the host returns media type `DK_UNKNOWN`. See *NOTES* for blocksize usage and capacity information.

DKIOCSMBOOT

The argument is a pointer to struct *mboot*.

Copies the *mboot* information supplied in the argument to the absolute sector 0 of the device. Prior to copying the information, this `ioctl()` performs the following checks on the *mboot* data:

- Ensures that the signature field is set to 0xAA55.
- Ensures that partitions do not overlap.

- On SPARC platforms, determines if the device is a removable media. If the above verification fails, `errno` is set to `EINVAL` and the `ioctl()` command fails.

x86 Platforms — Upon successful write of *mboot*, the partition map structure maintained in the driver is updated. If the new Solaris partition is different from the previous one, the internal VTOC table maintained in the driver is set as follows:

If `_SUNOS_VTOC_8` is defined:

Partition: 0. Start: 0. Capacity = Capacity of device.

Partition: 2. Start: 0. Capacity = Capacity of device.

If `_SUNOS_VTOC_16` is defined:

Partition: 2. Start: 0. Size = Size specified in *mboot* - 2 cylinders.

Partition: 8. Start: 0. Size = Sectors/cylinder.

Partition: 9. Start: Sectors/cylinder. Size = 2 * sectors/cylinder

To determine if the Solaris partition has changed:

If either offset or the size of the Solaris partition is different from the previous one then it shall be deemed to have changed. In all other cases, the internal VTOC info remains as before.

SPARC Platforms — The VTOC label and *mboot* both occupy the same location, namely sector 0. As a result, following the successful write of *mboot* info, the internal VTOC table maintained in the driver is set as follows:

Partition: 0. Start: 0. Size = Capacity of device.

Partition: 2. Start: 0. Size = Capacity of device.

See the NOTES section for usage of `DKIOCSMBOOT` when modifying Solaris partitions.

DKIOCGETVOLCAP

This `ioctl` provides information and status of available capabilities.

`vc_info` is a bitmap and the valid flag values are:

`DKV_ABR_CAP` - Capable of application-based recovery
`DKV_DMR_CAP` - Ability to read specific copy of data when multiple copies exist. For example, in a two way mirror, this `ioctl` is used to read each side of the mirror.

`vc_set` is a bitmap and the valid flag values are:

`DKV_ABR_CAP` - This flag is set if ABR has been set on a device that supports ABR functionality.
`DKV_DMR_CAP` - Directed read has been enabled.

Note: These capabilities are not required to be persistent across a system reboot and their persistence depends upon the implementation. For example, if the ABR capability for a DRL mirror simply clears the dirty-region list and subsequently stops updating this list, there is no reason for persistence because the VM recovery is a no-op. Conversely, if the ABR capability is applied to a non-DRL mirror to indicate that the VM should not perform a full recovery of the mirror following a system crash, the capability must be persistent so that the VM know whether or not to perform recovery.

Return Errors:

EINVAL — Invalid device for this operation.

ENOTSUP — Functionality that is attempted to be set is not supported.

DKIOCSETVOLCAP

This ioctl sets the available capabilities for the device. If a capability flag is not set in `vc_set`, that capability is cleared.

`vc_info` flags are ignored

`vc_set` valid flags are:

DKV_ABR_CAP - Flag to set application-based recovery. A device can successfully support ABR only if it is capable.
 DKV_DMR_CAP - Flag to set directed read.

```
int
ioctl(int , DKIODMR, vol_directed_rd *);
```

DKIODMR

This ioctl allows highly available applications to perform round-robin reads from the underlying devices of a replicated device.

```
vdr_offset      - offset at which the read should occur.
vdr_nbytes      - number of bytes to be read
vdr_bytesread   - number of bytes successfully read by the kernel.
vdr_data        - pointer to a user allocated buffer to return the data
                  read
vdr_side        - side to be read. Initialized to DKV_SIDE_INIT
vdr_side_name   - The volume name that has been read.
```

Valid `vdr_flags` are:

```
DKV_DMR_NEXT_SIDE (set by user)
DKV_DMR_DONE (return value)
DKV_DMR_ERROR (return value)
DKV_DMR_SUCCESS (return value)
DKV_DMR_SHORT (return value)
```

The calling sequence is as follows: The caller sets the `vdr_flags` to `DKV_DMR_NEXT_SIDE` and `vdr_side` to `DKV_SIDE_INIT` at the start. Subsequent calls should be made without any changes to these values. If they are changed the results of the ioctl are indeterminate.

When `DKV_SIDE_INIT` is set, the call results in the kernel reading from the first side. The kernel updates `vdr_side` to indicate the side that was read, and `vdr_side_name` to contain the name of that side. `vdr_data` contains the data that was read. Therefore to perform a round-robin read all of the valid sides, there is no need for the caller to change the contents of `vdr_side`.

Subsequent `ioctl` calls result in reads from the next valid side until all valid sides have been read. On success, the kernel sets `DKV_DMR_SUCCESS`. The following table shows the values of `vdr_flags` that are returned when an error occurs:

<code>vdr_flags</code>	<code>vdr_side</code>	Notes
<code>DKV_DMR_ERROR</code>	<code>DKV_SIDE_INIT</code>	No valid side to read
<code>DKV_DMR_DONE</code>	Not Init side	All valid sides read
<code>DKV_DMR_SHORT</code>	Any value	Bytes requested cannot be read. <code>vdr_bytesread</code> set to bytes actually read.

Typical code fragment:

```
enable->vc_set |= DKV_ABR_SET;
retval = ioctl(filedes, DKIOSETVOLCAP, enable);
if (retval != EINVAL || retval != ENOTSUP) {
    if (info->vc_set & DKV_DMR_SET) {
        dr->vdr_flags |= DKV_DMR_NEXT_SIDE;
        dr->vdr_side = DKV_SIDE_INIT;
        dr->vdr_nbytes = 1024;
        dr->vdr_offset = 0xff00;
        do {
            rval = ioctl(filedes, DKIODMR, dr);
            if (rval != EINVAL) {
                /* Process data */
            }
        } while (rval != EINVAL || dr->vdr_flags &
            (DKV_DMR_DONE | DKV_DMR_ERROR | DKV_DMR_SHORT))
    }
}
```

RETURN VALUES

Upon successful completion, the value returned is 0. Otherwise, -1 is returned and `errno` is set to indicate the error.

x86 Only

The following `ioctl()` requests set and/or retrieve the current disk controller, partitions, or geometry information on the x86 architecture.

<code>DKIOCG_PHYGEOM</code>	The argument is a pointer to a <code>dk_geom</code> structure (described below). This <code>ioctl()</code> gets the driver's notion of the physical geometry of the disk drive. It is functionally identical to the <code>DKIOCGGEOM</code> <code>ioctl()</code> .
<code>DKIOCG_VIRTGEOM</code>	The argument is a pointer to a <code>dk_geom</code> structure (described below). This <code>ioctl()</code> gets the controller's (and hence the driver's) notion of the virtual geometry of the disk drive. Virtual geometry is a view of the disk geometry maintained by the firmware in a host bus

adapter or disk controller. If the disk is larger than 8 Gbytes, this ioctl fails because a CHS-based geometry is not relevant or useful for this drive.

```

/*
 * Definition of a disk's geometry
 */
*/struct dk_geom {
    unsigned shor   dkg_ncyl;           /* # of data cylinders */
    unsigned shor   dkg_acyl;           /* # of alternate cylinders */
    unsigned short  dkg_bcyl;           /* cyl offset (for fixed head area) */
    unsigned short  dkg_nhead;          /* # of heads */
    unsigned short  dkg_obs1;           /* obsolete */
    unsigned short  dkg_nsect;          /* # of sectors per track*/
    unsigned short  dkg_intrlv;         /* interleave factor */
    unsigned short  dkg_obs2;           /* obsolete */
    unsigned short  dkg_obs3;           /* obsolete */
    unsigned short  dkg_apc;            /* alternates per cylinder */
                                           /* (SCSI only) */
    unsigned short  dkg_rpm;            /* revolutions per min*/
    unsigned short  dkg_p cyl;          /* # of physical cylinders */
    unsigned short  dkg_write_reinstruct; /* # sectors to skip, writes*/
    unsigned short  dkg_read_reinstruct; /* # sectors to skip, reads*/
    unsigned short  dkg_extra[7];       /* for compatible expansion*/
};

```

DKIOCADDBAD This `ioctl()` forces the driver to re-examine the alternates slice and rebuild the internal bad block map accordingly. It should be used whenever the alternates slice is changed by any method other than the `addbadsec(1M)` or `format(1M)` utilities. `DKIOCADDBAD` can only be used for software remapping on IDE drives; SCSI drives use hardware remapping of alternate sectors.

DKIOCPARTINFO The argument is a pointer to a `part_info` structure (described below). This `ioctl()` gets the driver's notion of the size and extent of the partition or slice indicated by the file descriptor argument.

```

/*
 * Used by applications to get partition or slice information
 */
struct part_info {
    daddr_t    p_start;
    int        p_length;
};

```

SEE ALSO `fdisk(1M)`, `format(1M)`, `ioctl(2)`, `sd(7D)`, `cdio(7I)`, `fdio(7I)`, `hdio(7I)`

x86 Only `addbadsec(1M)`, `cmdk(7D)`

NOTES | Blocksize information provided in `DKIOCGMEDIAINFO` is the size (in bytes) of the device's basic unit of operation and may differ from the blocksize that the Solaris operating environment exports to the user. Capacity information provided in the `DKIOCGMEDIAINFO` are for reference only and you are advised to use the values returned by `DKIOCGGEOM` or other appropriate `ioctl` for accessing data using the standard interfaces.

For x86 only: If the `DKIOCSMBOOT` command is used to modify the Solaris partitions, the VTOC information should also be set appropriately to reflect the changes to partition. Failure to do so will lead to unexpected results when the device is closed and re-opened fresh at a later time. This is because a default VTOC is assumed by driver when a Solaris partition is changed. The default VTOC will persist until the `ioctl` `DKIOCSVTOC` is called to modify VTOC or the device is closed and re-opened. At that point, the old valid VTOC will be read from the disk if it is still available.

dlcosmk(7ipp)

NAME	dlcosmk – Data Layer Class of Service Marker				
DESCRIPTION	<p>The <code>dlcosmk</code> marker is an action module that is executed as a result of classifying or metering packets. It marks the packet with a user priority defined by the <i>IEEE 801.D</i> standard. This feature is only possible on a VLAN device.</p> <p>The 3-bit user priority is part of the <i>802.1Q</i> VLAN header tag that is part of the ethernet header (carrying the IP packet).</p>				
STATISTICS	<p>The <code>dlcosmk</code> module exports the following statistics through <code>kstat</code>:</p> <p>Global statistics:</p> <pre>module: dlcosmk instance: <action id> name: dlcosmk statistics class <action name> crtime snaptime b_band <b_band value> dl_max <dl_max value> usr_pri <configured CoS> npackets <number of packets> epackets <number of packets in error> ipackets <number of packets not processed></pre>				
FILES	<code>/kernel/ipp/sparcv9/dlcosmk</code> 64-bit module (SPARC only.)				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWqos</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWqos
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWqos				
SEE ALSO	<code>ipqosconf(1M)</code> , <code>dscpmk(7IPP)</code> , <code>flowacct(7IPP)</code> , <code>ipqos(7IPP)</code> , <code>ipgpc(7IPP)</code> , <code>tokenmt(7IPP)</code> , <code>tswtclmt(7IPP)</code>				

NAME	dlpi – Data Link Provider Interface
SYNOPSIS	<pre>#include <sys/dlpi.h></pre>
DESCRIPTION	<p>SunOS STREAMS-based device drivers wishing to support the STREAMS TCP/IP and other STREAMS-based networking protocol suite implementations support Version 2 of the Data Link Provider Interface (DLPI). DLPI V2 enables a data link service user to access and use any of a variety of conforming data link service providers without special knowledge of the provider's protocol. Specifically, the interface is intended to support Ethernet, X.25 LAPB, SDLC, ISDN LAPD, CSMA/CD, FDDI, token ring, token bus, Bisync, and other datalink-level protocols.</p> <p>The interface specifies access to the data link service provider in the form of <code>M_PROTO</code> and <code>M_PCPROTO</code> type STREAMS messages and does not define a specific protocol implementation. The interface defines the syntax and semantics of primitives exchanged between the data link user and the data link provider to attach a physical device with physical-level address to a stream, bind a datalink-level address to the stream, get implementation-specific information from the data link provider, exchange data with a peer data link user in one of three communication modes (connection, connectionless, acknowledged connectionless), enable/disable multicast group and promiscuous mode reception of datalink frames, get and set the physical address associated with a stream, and several other operations.</p> <p>Solaris conforms to The Open Group Technical Standard for <i>DLPI, Version 2</i>. For free access to this specification, point your browser to www.opengroup.org/pubs/catalog/c811.htm. Solaris also provides extensions to the DLPI standard, as detailed in this man page.</p>
SOLARIS-SPECIFIC DLPI EXTENSIONS	<p>Notification Support</p> <p>Enables DLPI consumers to register for notification when events of interest occur at the DLPI provider. The negotiation may be performed on any attached DLPI stream, and begins with the DLPI consumer, sending a <code>DL_NOTIFY_REQ</code> to the provider, which is an <code>M_PROTO</code> message with the following payload:</p> <pre>typedef struct { t_uscalar_t dl_primitive; uint32_t dl_notifications; uint32_t dl_timelimit; } dl_notify_req_t;</pre> <p>The <code>dl_primitive</code> field must be set to <code>DL_NOTIFY_REQ</code>; the <code>dl_timelimit</code> field is reserved for future use and must be set to zero. The <code>dl_notifications</code> field is a bitmask containing the event types the consumer is interested in receiving, and must be zero or more of:</p> <pre>DL_NOTE_LINK_DOWN /* notify when link has gone down */ DL_NOTE_LINK_UP /* notify when link has come up */ DL_NOTE_SDU_SIZE /* notify when link mtu changes */ DL_NOTE_SPEED /* notify when link speed changes */</pre> <p>Consumers may find it useful to send a <code>DL_NOTIFY_REQ</code> message with no requested types to check if the DLPI provider supports the extension.</p>

dlpi(7P)

Upon receiving the `DL_NOTIFY_REQ`, the DLPI provider must generate a `DL_NOTIFY_ACK`, which is an `M_PROTO` message with the following payload:

```
typedef struct {
    t_uscalar_t    dl_primitive;
    uint32_t       dl_notifications;
} dl_notify_ack_t;
```

The `dl_primitive` field must be set to `DL_NOTIFY_ACK`. The `dl_notifications` field must include any requested notifications that the driver supports, along with any other unrequested notifications that the driver supports. However, regardless of the notifications the driver supports, it is restricted to sending only `DL_NOTIFY_IND` messages (see below) that were requested in the `DL_NOTIFY_REQ`.

Since there are additional notification types which are not yet available for public use, DLPI consumers and providers must take care when inspecting and setting the `dl_notifications` field. Specifically, consumers must be careful to only request the above notification types, and providers must be careful to not include any unrecognized notification types in the `dl_notifications` field when constructing the `DL_NOTIFY_ACK`.

DLPI consumers may receive a `DL_ERROR_ACK` message (with `dl_primitive` set to `DL_NOTIFY_REQ`) in response to the initial `DL_NOTIFY_REQ` message. This message indicates that the DLPI provider does not support the DLPI notification extension. Otherwise, the DLPI consumer will receive a `DL_NOTIFY_ACK` and should expect to receive `DL_NOTIFY_IND` messages for any types that it requested that were still set in it. The `DL_NOTIFY_IND` is an `M_PROTO` message with the following payload:

```
typedef struct {
    t_uscalar_t    dl_primitive;
    uint32_t       dl_notification;
    uint32_t       dl_data;
    t_uscalar_t    dl_addr_length;
    t_uscalar_t    dl_addr_offset;
} dl_notify_ind_t;
```

The `dl_primitive` field must be set to `DL_NOTIFY_IND`, and the `dl_notification` field must be set to the event type that has occurred (e.g., `DL_NOTE_LINK_DOWN`). only a single event type may be set in each `DL_NOTIFY_IND`. For the `DL_NOTE_SPEED` event type, `dl_data` must be set to the current interface speed in kilobits per second, and for the `DL_NOTE_SDU_SIZE` event type, `dl_data` must be set to the current MTU in bytes; otherwise, `dl_data` must be set to zero. The other members of this structure are not yet available for public use and must be set to zero by DLPI providers and ignored by DLPI consumers.

In addition to generating `DL_NOTIFY_IND` messages when a requested event has occurred, the DLPI provider must initially generate one or more `DL_NOTIFY_IND` messages to notify the DLPI consumer of the the current state of the interface. For instance, if the consumer has requested `DL_NOTE_LINK_UP` |

DL_NOTE_LINK_DOWN, the provider must send a DL_NOTIFY_IND containing the current state of the link (either DL_NOTE_LINK_UP or DL_NOTE_LINK_DOWN) after sending the DL_NOTIFY_ACK.

The DLPI provider must continue to generate DL_NOTIFY_IND messages until it receives a new DL_NOTIFY_REQ message or the DLPI stream is detached (or closed). Further, a DLPI style 2 driver must keep track of the requested events after a DL_DETACH_REQ operation, and if a subsequent DL_ATTACH_REQ is received, it must send gratuitous DL_NOTIFY_IND messages to notify the consumer of the current state of the device, since the state may have changed while detached (or the consumer may have simply discarded its previous state).

FILES Files in or under /dev.

SEE ALSO hme(7D), ge(7D), qfe(7D), glid(7D)

NOTES Streams drivers for network interface cards (NIC) must meet the following driver name constraints:

- **Length** — Name cannot exceed 16 characters. Names containing three to eight characters are preferred.
- **Legal Characters** — Legal characters are: alphanumeric (a-z, A-Z, 0-9), and the underscore ('_'). Additionally, the first and/or last character of a name cannot be a digit.

dmfe(7D)

NAME	dmfe – Davicom Fast Ethernet driver for Davicom DM9102A
SYNOPSIS	<code>/platform/sun4u/kernel/drv/sparcv9/dmfe</code>
DESCRIPTION	<p>The dmfe Ethernet device provides 100Base-TX networking interfaces using the Davicom DM9102A chip, which incorporates its own internal transceiver.</p> <p>The dmfe driver is a multithreaded, loadable, clonable, GLD-based STREAMS driver. Multiple controllers installed within the system are supported by the driver. The dmfe driver functions include controller initialization, frame transmit and receive, promiscuous and multicast support, and error recovery and reporting.</p> <p>The 100Base-TX standard specifies an auto-negotiation protocol to automatically select the mode and speed of operation. The internal transceiver is capable of performing autonegotiation with the remote-end of the link (link partner) and receives the capabilities of the remote end. It selects the highest common denominator mode of operation based on the priorities. The internal transceiver also supports a forced-mode of operation under which the driver selects the operational mode.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The <code>/dev/dmfe</code> cloning character-special device is used to access all Davicom DM9102A devices installed in the system.</p> <p>The dmfe driver is dependent on <code>/kernel/misc/gld</code>, a loadable kernel module that provides the dmfe driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on the primitives supported by the driver.</p> <p>You must send an explicit <code>DL_ATTACH_REQ</code> message to associate the opened stream with a particular device (ppa). The ppa ID is interpreted as an unsigned integer data type and indicates the corresponding device instance (unit) number. If the ppa field value does not correspond to a valid device instance number for this system, an error (<code>DL_ERROR_ACK</code>) is returned. The device is initialized on first attach and de-initialized (stopped) at last detach.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to a <code>DL_INFO_REQ</code> are as follows:</p> <ul style="list-style-type: none">■ Maximum SDU is 1500 (<code>ETHERMTU</code> - defined in <code>sys/ethernet.h</code>).■ Minimum SDU is 0.■ DLSAP address length is 8.■ MAC type is <code>DL_ETHER</code>.■ The sap length value is -2, meaning the physical address component is followed immediately by a 2-byte sap component within the DLSAP address.■ The broadcast address value is the Ethernet/IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>). <p>Once in the <code>DL_ATTACHED</code> state, you must send a <code>DL_BIND_REQ</code> to associate a particular Service Access Point (SAP) with the stream.</p>

CONFIGURATION By default, the `dmfe` driver performs auto-negotiation to select the speed and mode of the link. Link speed and mode can be 100 Mbps (full or half-duplex) or 10 Mbps (full or half-duplex) as described in the 100Base-TX standard.

The auto-negotiation protocol automatically selects speed mode (either 100 Mbps or 10 Mbps) and operation mode (either full-duplex or half-duplex) as the highest common denominator supported by both link partners. Because the `dmfe` device supports all modes, this effectively selects the highest-throughput mode supported by the other device.

Alternatively, you can explicitly specify the link parameters by adding entries to the `dmfe` driver configuration file (`/platform/sun4u/kernel/drv/dmfe.conf`). You can set the `speed` parameter to 10 or 100 to force `dmfe` devices to operate at the specified speed. Additionally, you can set the `full-duplex` parameter to 0 or 1 to disable or force full-duplex operation, respectively.

Note that specifying either "speed" or "full-duplex" explicitly disables auto-negotiation. To enable the driver to determine the appropriate setting for each parameter, you should always set both parameters. If it is necessary to force either speed or duplex setting (for example, because the `dmfe` device is connected to an ancient device or hub that does not support auto-negotiation), both parameters should be explicitly specified to match the requirements of the external device.

FILES `/dev/dmfe`
Character special device

`/platform/sun4u/kernel/drv/dmfe.conf`
dmfe configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`, `gld(7D)`, `dlpi(7P)`, `streamio(7I)`

Writing Device Drivers

STREAMS Programming Guide

Network Interfaces Programmer's Guide

dnet(7D)

NAME	dnet – Ethernet driver for DEC 21040, 21041, 21140 Ethernet cards
SYNOPSIS	<code>/kernel/drv/dnet</code>
DESCRIPTION	<p>The dnet Ethernet driver is a multithreaded, loadable, clonable, STREAMS GLD driver. Multiple controllers installed within the system are supported by the driver. The dnet driver functions include controller initialization, frame transmit and receive, functional addresses, promiscuous and multicast support, and error recovery and reporting.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device, <code>/dev/dnet</code>, is used to access all DEC 21040/21041/21140 devices installed in the system.</p> <p>The dnet driver is dependent on <code>/kernel/misc/gld</code>, a loadable kernel module that provides the dnet driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on the primitives supported by the driver.</p> <p>The device is initialized on the first attach and de-initialized (stopped) on the last detach.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to a <code>DL_INFO_REQ</code> from the user are as follows:</p> <ul style="list-style-type: none">■ The maximum SDU is 1500 (<code>ETHERMTU</code> - defined in <code><sys/ethernet.h></code>).■ The minimum SDU is 0.■ The DLSAP address length is 8.■ The MAC type is <code>DL_ETHER</code>.■ The sap length value is -2, meaning the physical address component is followed immediately by a 2-byte sap component within the DLSAP address.■ The broadcast address value is the Ethernet/IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>). <p>Once in the <code>DL_ATTACHED</code> state, the user must send a <code>DL_BIND_REQ</code> to associate a particular Service Access Point (SAP) with the stream.</p>
PRECONFIGURATION	The PCI configuration process varies from system to system. Follow the instructions provided by the vendor.
Known Problems and Limitations	<ul style="list-style-type: none">■ On multiport cards (exception: Osicom (Rockwell) RNS2340), the first port is the top port. (On the Osicom RNS2340, the first port is the bottom port.)■ If the dnet driver fails to determine the correct speed and duplex mode resulting in a corresponding drop in performance, set the speed and duplex mode using the <code>dnet.conf</code> file.■ The dnet driver incorrectly counts carrier lost or no carrier errors while in full-duplex mode. There is no carrier signal present when in full-duplex mode and it should not be counted as an error.■ Version 4 SRAM formats are not supported.
CONFIGURATION	The <code>/kernel/drv/dnet.conf</code> file supports the following options:

`full-duplex` For full duplex operation use `full-duplex=1`, for half duplex use `full-duplex=0`. Half-duplex operation gives better results on older 10mbit networks.

`speed` For 10mbit operation use `speed=10`, for 100mbit operation use `speed=100`. Certain 21140 based cards will operate at either speed. Use the `speed` property to override the 100mbit default in this case.

FILES `/dev/dnet` character special device
`/kernel/drv/dnet.conf` dnet configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`, `dlpi(7P)`, `gld(7D)` `streamio(7I)`

Writing Device Drivers

STREAMS Programming Guide

Network Interfaces Programmer's Guide

dpt(7D)

NAME	dpt – DPT ServeRAID IV SCSI host bus adapter and RAID adapter driver	
DESCRIPTION	<p>The <code>dpt</code> driver is a driver for the DPT (Distributed Processing Technology) family of SmartRAID IV SCSI HBA and RAID adapters. The following HBA adapters are supported: PM2024, PM2044UW, PM2044W, PM2124, PM2124W, PM2144UW, and PM2144W.</p> <p>The following RAID adapters are supported: PM3224, PM3224W, PM3334UW, and PM3334W.</p>	
PRECONFIGURATION	<p>ON DPT PM3224 <i>only</i>: Only EPROM 7A and later versions are supported.</p> <ul style="list-style-type: none">■ DPT PM2024 and PM2124 <i>only</i>: Only EPROM 6D4 and later versions are supported.■ Use adapters with SmartROM version 3.B or later versions only.■ Be sure that the controller board is installed in a PCI bus-mastering slot.■ Disable PCI parity checking if your firmware version is earlier than version 7A, if your system memory is ECC, or if your system does not check parity.	
Known Problems and Limitations	During system boot, a message may be displayed saying a DPT controller driver cannot be installed. This message indicates that the motherboard installed in your system may contain ECC memory or may not check parity. If you see this message is displayed, disable PCI parity checking.	
Supported Settings	<ul style="list-style-type: none">■ I/O Address: Auto	
CONFIGURATION	<p>Auto-configuration code determines whether the adapter is present at the configured address and what types of devices are attached to it. The DPT ServeRAID is primarily used as a disk array (system drive) controller.</p> <p>To configure the attached disk arrays, you must configure the controller (using the configuration utilities provided by the hardware manufacturer) before you boot the Solaris operating environment. You use the configuration utilities to set RAID levels, stripe parameters, cache mechanisms and perform other functions. For more information, see the user manual supplied with your hardware.</p>	
FILES	<code>/kernel/drv/dpt.conf</code>	dpt configuration file
	<code>/dev/dsk/cndn [s p] n</code>	block device
	<code>/dev/rdisk/cndn [s p] n</code>	raw device where:
	<code>cn</code>	controller <i>n</i>
	<code>dn</code>	LUN <i>n</i> (0–7)
	<code>sn</code>	UNIX system slice <i>n</i> (0–15)
	<code>pn</code>	<code>fdisk(1M)</code> partition (0)

dpt(7D)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `fdisk(1M)`, `attributes(5)`, `cmdk(7D)`

dr(7d)

NAME	dr, drmach, ngdr, ngdrmach – Sun Enterprise 10000 dynamic reconfiguration driver
SYNOPSIS	dr drmach ngdr ngdrmach
DESCRIPTION	<p>The dynamic reconfiguration (DR) driver consists of a platform-independent driver and a platform-specific module. The DR driver uses standard features of the Solaris operating environment whenever possible to control DR operations and calls the platform specific module as needed. The DR driver creates minor nodes in the file system that serve as attachment points for DR operations.</p> <p>The DR driver provides a pseudo-driver interface to sequence attach and detach operations on system boards using file system entry points referred to as "attachment points." The attachment point form depends on the platform.</p>
Sun Enterprise 10000 Server	<p>On the Sun Enterprise 10000 server, the DR driver consists of a platform-independent driver (ngdr) and a platform-specific module (ngdrmach).</p> <p>The domain configuration server (DCS) accepts DR requests from the system services processor (SSP) and uses the <code>libcfgadm(3LIB)</code> interface to initiate the DR operation. After the operation is performed, the results are returned to the SSP. For more information about the DCS on the Sun Enterprise 10000, refer to the <code>dcS(1M)</code> man page and the <i>Sun Enterprise 10000 Dynamic Reconfiguration User Guide</i>.</p> <p>The DR driver creates physical attachment points for system board slots that takes the following form:</p> <pre>/devices/pseudo/ngdr@0:SBx</pre> <p>Where <i>x</i> represents the slot number (0 to 15) for a particular board.</p> <p>The <code>cfgadm_sbd(1M)</code> plugin creates dynamic attachment points that refer to components on system boards, including CPUs, memory, or I/O devices. Refer to the <code>cfgadm_sbd(1M)</code> man page for more details.</p>
SEE ALSO	<code>cfgadm_sbd(1M)</code> , <code>ioctl(2)</code> , <code>libcfgadm(3LIB)</code> <i>Sun Enterprise 10000 Dynamic Reconfiguration User Guide</i>

NAME dscpmk – Differentiated Services Code Point Marker

DESCRIPTION The dscpmk marker is an action module that is executed as a result of classifying or metering packets. It sets a codepoint in the IP header as defined in *RFC-2474: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*.

STATISTICS The dscpmk module exports the following statistics available through `kstat`:

Global statistics:

```

module: dscpmk                instance: <action id>
  name: dscpmk stats          class <action name>
  crtime
  snaptime
  npackets                    <number of packets>
  epackets                    <number of packets in error>
  ipackets                    <number of packets not processed>
  dscp_unchanged              <number of packets with DSCP unchanged>
  dscp_changed                <number of packets with DSCP changed>

```

Also, for each DSCP the following is exported:

```

module: dscpmk                instance: <action id>
  name: dscpmk_dscp0x<DSCP> value class: <action name>
  dscp                        <DSCP value>
  npackets                    <number of packets for this DSCP>

```

FILES /kernel/ipp/sparcv9/dscpmk
64-bit module (SPARC only.)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWqos

SEE ALSO `ipqosconf(1M)`, `dlcosmk(7IPP)`, `flowacct(7IPP)`, `ipqos(7IPP)`, `ipgpc(7IPP)`, `tokenmt(7IPP)`, `tswtclmt(7IPP)`

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers K. Nichols, S. Blake, F. Baker, D. Black — The Internet Society, 1998

dtrace(7D)

NAME dtrace – DTrace dynamic tracing facility

DESCRIPTION The `dtrace` driver provides the dynamic instrumentation and tracing facilities for the DTrace software, as well as the built-in `dtrace` provider. The `dtrace` driver is not a public interface and you access the instrumentation offered by this provider through DTrace tools such as `dtrace(1M)`. Refer to the *Solaris Dynamic Tracing Guide* for a description of the public documented interfaces available for the DTrace facility and `dtrace` provider probes.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdtrp
Interface Stability	Private

SEE ALSO `dtrace(1M)`, `libdtrace(3LIB)`, `attributes(5)`

Solaris Dynamic Tracing Guide

NAME	e1000g – 1000g Intel Gigabit and 82546EB based network interface controllers																		
SYNOPSIS	<code>/dev/e1000g</code>																		
DESCRIPTION	The e1000g Ethernet driver is a multithreaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code> , over Intel Gigabit controllers. Multiple Intel Gigabit controllers installed within the system are supported by the driver. The e1000g driver provides basic support for the Intel Gigabit controllers. Functions include chip initialization, frame transmit and receive, multicast support, and error recovery and reporting.																		
APPLICATION PROGRAMMING INTERFACE	The cloning, character-special device <code>/dev/e1000g</code> is used to access all Intel Gigabit devices installed within the system.																		
e1000g and DLPI	The e1000g driver is a self-sufficient DLPI module that provides all the LAN driver and STREAMS functionality; the <code>/kernel/misc/gld</code> module is not required.																		
	The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> from the user are as follows: <ul style="list-style-type: none"> ■ The maximum SDU (with jumbo frame) is 16128 (ETHERMTU). ■ The minimum SDU is 0. The driver pads to the mandatory 60-octet minimum packet size. ■ The <code>dlsap</code> address length is 8. ■ The MAC type is <code>DL_ETHER</code>. ■ The <code>sap</code> length value is -2, meaning the physical address component is followed immediately by a 2-byte <code>sap</code> component within the DLSAP address. ■ The broadcast address value is Ethernet/IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>). 																		
CONFIGURATION	The e1000g driver does not support the use of shared RAM on the board.																		
	To configure the e1000g driver: <ul style="list-style-type: none"> ■ Use <code>prtconf -v grep pci8086, [12] [01] [01] [0-F]</code> to obtain the instance number of the driver. ■ Use <code>ifconfig e1000ginstance plumb</code> to plumb the controller. ■ Use <code>ifconfig e1000ginstance inet ip_address netmask + broadcast + -trailers up</code> to bring up the interface. ■ Use the <code>ping(1M)</code> command to contact interfaces on the network to verify that the configuration is operational. 																		
Configuration File Options	The following <code>e100g.conf</code> configuration options are supported: <p>AutoNegAdvertised</p> <p>This is a bitmap for the speeds advertised during autonegotiation.</p> <table border="1"> <tr> <td>Bit</td> <td>7</td> <td>6</td> <td>5</td> <td>4</td> <td>3</td> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>Setting</td> <td>N/A</td> <td>N/A</td> <td>1000F</td> <td>N/A</td> <td>100F</td> <td>100H</td> <td>10F</td> <td>10H</td> </tr> </table>	Bit	7	6	5	4	3	2	1	0	Setting	N/A	N/A	1000F	N/A	100F	100H	10F	10H
Bit	7	6	5	4	3	2	1	0											
Setting	N/A	N/A	1000F	N/A	100F	100H	10F	10H											

e1000g(7D)

The adapter only autonegotiates to a speed which is advertised. For example: `AutoNegAdvertised = 4` causes an adapter to only advertise autonegotiation at 100 Mbps, full duplex. No other link speeds is accepted or given during autonegotiation. `AutoNegAdvertised=47` advertises all speeds available, This is the same as using the default setting of 0.

0-255 Allowed values

0 Default

DmaFairness

When enabled, direct memory access by the driver is fair, meaning transmit and receive have equal priority. When disabled, receive takes priority over transmit.

The allowed values are:

0 Disable.

1 Enable. This is the default.

FlowControl

Flow control utilizes ethernet XON and unicast and multicast XOFF packets to allow ethernet equipment to slow down the stream of data between two ethernet devices.

The allowed values are:

0 Disable. Packets can get dropped in high throughput situations, leading to reduced network performance.

1 Receive only.

2 Transmit only.

3 Receive and transmit. This is the default.

4 Use adapter's EEPROM-programmed factory default setting.

ForceSpeedDuplex

Specify the speed and duplex mode for each instance.

If you set `ForceSpeedDuplex=7, 4`, the `e1000g0` is set to autonegotiate and `e1000g1` is set to 100 Mbps, full duplex.

Fiber optic ethernet adapters ignore this setting.

The allowed values are:

1 10 Megabits per second, Half Duplex.

2 10 Megabits per second, Full Duplex.

3 100 Megabits per second, Half Duplex.

4 100 Megabits per second, Full Duplex.

7 Autonegotiate speed and duplex. This is the default.

MaxFrag

Maximum number of STREAMS memory buffer fragmentation of incoming or outgoing packets. This value can effect the performance and memory requirements of the driver.

2-8 Allowed values

8 Default

MaxFrameSize

Upper limit on the maximum MTU size the driver allows. All Intel gigabit adapters, except the 82542-based Intel PRO/1000 adapter, allow the configuration of jumbo frames.

The maximum MTU accepted by the MAC is 16128. Use `ifconfig(1M)` To configure jumbo frames. `ifconfig` with the adapter instance and the `mtu` argument (`ifconfig e1000g0 mtu 16128`) configures adapter `e1000g0` for the maximum allowable jumbo frame size.

The allowed values are:

0 Standard ethernet frames with a MTU equal to 1500. This is the default.

1 Jumbo frames with a maximum MTU of 4096.

2 Jumbo frames with a maximum MTU of 8192.

3 Jumbo frames with a maximum MTU of 16384.

MaxFreeListBuf

Maximum number or pre-allocated buffers the driver can use for received data. Pre-allocating buffers can help with receive performance. It also increases the memory requirements of the driver.

0-4096 Allowed values

256 Default

MaxNumReceivePackets

Maximum number of receive packets that the driver can handle for each interrupt.

CPU utilization can be lowered through more efficient interrupt management. If this value is increased, the time needed by the CPU to process the individual interrupts can increase and nullify any performance gains realized by handling less interrupts.

16-1024 Allowed values

32 Default

MWIEnable

x86 Architecture servers based on the Intel 450NX chipset are known to have problems with Memory Write and Invalidate. Disable this feature only for this chipset.

The allowed values are:

e1000g(7D)

- 0 Disable. This is the default.
- 1 Enable.

NumRxDescriptors

Number of receive descriptors available to the driver. More receive descriptors increase receive performance at the penalty of decreased available memory.

If you set this value to 256 on Solaris 7 when multiple Intel Gigabit NICs are configured, it causes memory allocation problems. It is recommended that you set this value lower when using multiple Intel Gigabit adapters on Solaris 7.

- 80–256 Allowed values
- 256 Default

NumTxDescriptors

Number of transmit descriptors available to the driver. More transmit descriptors increase transmit performance at the penalty of decreased available memory.

If you set this value to 256 on Solaris 7 when multiple Intel Gigabit NICs are configured, it causes memory allocation problems. It is recommended that you set this value lower when using multiple Intel Gigabit adapters on Solaris 7.

- 80–256 Allowed values
- 256 Default

ProfileJumboTraffic

If you need to collect statistics on jumbo frame usage, enabling this parameter causes the `e1000g` driver to keep a count of the number of jumbo frames transmitted and received. Use the `netstat -k` command to view jumbo frame utilization statistics. See `netstat(1M)`

The allowed values are:

- 0 Disable. This is the default.
- 1 Enable.

TbiCompatibilityEnable=1;

You must enable this feature on Intel 82543CG-based copper adapters to operate correctly with TBI mode ethernet hardware.

The allowed values are:

- 0 Disable.
- 1 Enable. This is the default.

TxInterruptDelay

Amount of time (in 1.024 micro second units) between the time transmit data is queued in a transmit descriptor and the transmit interrupt is sent.

- 0–65535 Allowed values. A value of 0 completely disables any transmit interrupt delay. Some delay is beneficial in reducing CPU utilization by the driver.

300 Default.

FILES

/dev/e1000g	
e1000g	Character special device
/kernel/drv/e1000g.conf	Configuration file for 32-bit driver
/kernel/drv/amd64/e1000g.conf	Configuration file for 64-bit driver
sys/stropts.h	stropts network header file
sys/ethernet.h	Ethernet network header file
sys/dlpi.h	dlpi network header file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `ifconfig(1M)`, `netstat(1M)`, `ping(1M)`, `attributes(5)`, `dlpi(7P)`

Intel PRO/1000 Gigabit Adapter Driver Installation Notes for Solaris

ecpp(7D)

NAME	ecpp – IEEE 1284 compliant parallel port driver
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ecppio.h> ecpp@unit-address</pre>
DESCRIPTION	<p>The <code>ecpp</code> driver provides a bi-directional interface to <i>IEEE 1284</i> compliant devices as well as a forward single-directional interface to Centronics devices. In addition to the Centronics protocol, the <code>ecpp</code> driver supports the <i>IEEE 1284</i> Compatibility, Nibble, and ECP protocols. <code>ECPP_COMPAT_MODE</code> and <code>ECPP_CENTRONICS</code> modes of operation have logically identical handshaking protocols, however devices that support <code>ECPP_COMPAT_MODE</code> are <i>IEEE 1284</i> compliant devices. <i>IEEE 1284</i> compliant devices support at least <code>ECPP_COMPAT_MODE</code> and <code>ECPP_NIBBLE_MODE</code>. Centronics devices support only <code>ECPP_CENTRONICS</code> mode.</p> <p>By default, <code>ECPP_COMPAT_MODE</code> devices have a strobe handshaking pulse width of 500ns. For this mode, forward data transfers are conducted by DMA. By default, the strobe pulse width for <code>ECPP_CENTRONICS</code> devices is two microseconds. Forward transfers for these devices are managed through PIO. The default characteristics for both <code>ECPP_COMPAT_MODE</code> and <code>ECPP_CENTRONICS</code> devices may be changed through tunable variables defined in <code>ecpp.conf</code>.</p> <p>The <code>ecpp</code> driver is an <i>exclusive-use</i> device, meaning that if the device is already open, subsequent opens fail with <code>EBUSY</code>.</p>
Default Operation	<p>Each time the <code>ecpp</code> device is opened, the device is marked as <code>EBUSY</code> and the configuration variables are set to their default values. The <code>write_timeout</code> period is set to 90 seconds.</p> <p>The driver sets the mode variable according to the following algorithm: The driver initially attempts to negotiate the link into <code>ECPP_ECP_MODE</code> during <code>open(2)</code>. If it fails, the driver tries to negotiate into <code>ECPP_NIBBLE_MODE</code> mode. If that fails, the driver operates in <code>ECPP_CENTRONICS</code> mode. Upon successfully opening the device, <i>IEEE 1284</i> compliant devices will be left idle in either reverse idle phase of <code>ECPP_ECP_MODE</code> or in <code>ECPP_NIBBLE_MODE</code>. Subsequent calls to <code>write(2)</code> invokes the driver to move the link into either <code>ECPP_COMPAT_MODE</code> or the forward phase of <code>ECPP_ECP_MODE</code>. After the transfer completes, the link returns to idle state.</p> <p>The application may attempt to negotiate the device into a specific mode or set the <code>write_timeout</code> values through the <code>ECPPIOC_SETPARMS ioctl(2)</code> call. For mode negotiation to be successful, both the host workstation and the peripheral must support the requested mode.</p>
Tunables	<p>Characteristics of the <code>ecpp</code> driver may be tuned by the variables described in <code>/kernel/drv/ecpp.conf</code>. These variables are read by the kernel during system startup. To tune the variables, edit the <code>ecpp.conf</code> file and invoke <code>update_drv(1M)</code> to have the kernel read the file again.</p>

	Some Centronics peripherals and certain <i>IEEE 1284</i> compatible peripherals will not operate with the parallel port operating in a fast handshaking mode. If printing problems occur, set "fast-centronics" and "fast-1284-compatible" to "false." See <code>/kernel/drv/ecpp.conf</code> for more information.
Read/Write Operation	The <code>ecpp</code> driver is a full duplex STREAMS device driver. While an application is writing to an <i>IEEE 1284</i> compliant device, another thread may read from it.
Write Operation	A <code>write(2)</code> operation returns the number of bytes successfully written to the stream head. If a failure occurs while a Centronics device is transferring data, the content of the status bits will be captured at the time of the error and can be retrieved by the application program using the <code>BPPIOC_GETERR</code> <code>ioctl(2)</code> call. The captured status information is overwritten each time an attempted transfer or a <code>BPPIOC_TESTIO</code> <code>ioctl(2)</code> occurs.
Read Operation	If a failure or error condition occurs during a <code>read(2)</code> , the number of bytes successfully read is returned (short read). When attempting to read a port that has no data currently available, <code>read(2)</code> returns 0 if <code>O_NDELAY</code> is set. If <code>O_NONBLOCK</code> is set, <code>read(2)</code> returns -1 and sets <code>errno</code> to <code>EAGAIN</code> . If <code>O_NDELAY</code> and <code>O_NONBLOCK</code> are clear, <code>read(2)</code> blocks until data become available.
IOCTLS	<p>The <code>ioctl(2)</code> calls described below are supported. Note that when <code>ecpp</code> is transferring data, the driver waits until the data has been sent to the device before processing the <code>ioctl(2)</code> call.</p> <p>The <code>ecpp</code> driver supports <code>prnio(7I)</code> interfaces.</p> <p>Note – The <code>PRNIOC_RESET</code> command toggles the <code>nInit</code> signal for 2 ms, followed by default negotiation.</p> <p>The following <code>ioctl(2)</code> calls are supported for backward compatibility and are not recommended for new applications:</p> <p>ECPPIOC_GETPARMS Get current transfer parameters. The argument is a pointer to a struct <code>ecpp_transfer_parms</code>. See below for a description of the elements of this structure. If no parameters have been configured since the device was opened, the structure will be set to its default configuration. See Default Operation above for more information.</p> <p>ECPPIOC_SETPARMS Set transfer parameters. The argument is a pointer to a struct <code>ecpp_transfer_parms</code>. If a parameter is out of range, <code>EINVAL</code> is returned. If the peripheral or host device cannot support the requested mode, <code>EPROTONOSUPPORT</code> is returned. See below for a description of <code>ecpp_transfer_parms</code> and its valid parameters.</p> <p>The Transfer Parameters Structure is defined in <code><sys/ecppio.h></code>.</p> <pre>struct ecpp_transfer_parms { int write_timeout;</pre>

ecpp(7D)

```
int mode;
};
```

The `write_timeout` field is set to the value of `ecpp-transfer-timeout` specified in the `ecpp.conf`. The `write_timeout` field specifies how long the driver will wait for the peripheral to respond to a transfer request. The value must be greater than 0 and less than `ECPP_MAX_TIMEOUT`. All other values are out of range.

The `mode` field reflects the *IEEE 1284* mode to which the parallel port is currently configured. The mode may be set to one of the following values only: `ECPP_CENTRONICS`, `ECPP_COMPAT_MODE`, `ECPP_NIBBLE_MODE`, `ECPP_ECP_MODE`. All other values are invalid. If the requested mode is not supported, `ECPP_IOC_SETPARMS` will return `EPROTONOSUPPORT` and the mode will be set to `ECPP_CENTRONICS` mode. Afterwards, the application may change the mode back to the original mode with `ECPP_IOC_SETPARMS`.

ECPP_IOC_GETDEVID

This ioctl gets the *IEEE 1284* device ID from the peripheral in specified mode. Currently, the device ID can be retrieved only in Nibble mode. A pointer to the structure defined in `<sys/ecppsys.h>` must be passed as an argument.

The 1284 device ID structure:

```
struct ecpp_device_id {
    int mode; /* mode to use for reading device id */
    int len; /* length of buffer */
    int rlen; /* actual length of device id string */
    char *addr; /* buffer address */
};
```

The mode is the *IEEE 1284* mode into which the port will be negotiated to retrieve device ID information. If the peripheral or host do not support the mode, `EPROTONOSUPPORT` is returned. Applications should set mode to `ECPP_NIBBLE_MODE`. `len` is the length of the buffer pointed to by `addr`. `rlen` is the actual length of the device ID string returned from the peripheral. If the returned `rlen` is greater than `len`, the application can call `ECPP_IOC_GETDEVID` again with a buffer length equal or greater than `rlen`. Note that the two length bytes of the *IEEE 1284* device ID are not taken into account and are not returned in the user buffer.

After `ECPP_IOC_GETDEVID` successfully completes, the driver returns the link to `ECPP_COMPAT_MODE`. The application is responsible for determining the previous mode the link was operating in and returning the link to that mode.

ECPP_IOC_TESTIO

Tests the forward transfer readiness of a peripheral operating in Centronics or Compatibility mode.

`TESTIO` determines if the peripheral is ready to receive data by checking the open flags and the Centronics status signals. If the current mode of the device is `ECPP_NIBBLE_MODE`, the driver negotiates the link into `ECPP_COMPAT_MODE`, check the status signals and then return the link to `ECPP_NIBBLE_MODE` mode. If

the current mode is `ECPP_CENTRONICS` or `ECPP_COMPAT_MODE`, `TESTIO` examines the Centronics status signals in the current mode. To receive data, the device must have the `nErr` and `Select` signals asserted and must not have the `PE` and `Busy` signals asserted. If `ecpp` is transferring data, `TESTIO` waits until the previous data sent to the driver is delivered before executing `TESTIO`. However if an error condition occurs while a `TESTIO` is waiting, `TESTIO` returns immediately. If `TESTIO` determines that the conditions are ok, 0 is returned. Otherwise, -1 is returned, `errno` is set to `EIO` and the state of the status pins is captured. The captured status can be retrieved using the `BPPIOC_GETERR` `ioctl(2)` call. The `timeout_occurred` and `bus_error` fields will never be set by this `ioctl(2)`. `BPPIOC_TESTIO` and `BPPIOC_GETERR` are compatible to the `ioctls` specified in [bpp\(7D\)](#).

BPPIOC_GETERR

Get last error status. The argument is a pointer to a `struct bpp_error_status` defined in `<sys/bpp_io.h>` header file. The error status structure is:

```
struct bpp_error_status {
    char    timeout_occurred; /* 1=timeout */
    char    bus_error;        /* not used */
    uchar_t pin_status;      /* status of pins which
                             /* could cause error */
};
```

The `pin_status` field indicates possible error conditions. The valid bits for `pin_status` are: `BPP_ERR_ERR`, `BPP_SLCT_ERR`, `BPP_PE_ERR`, `BPP_BUSY_ERR`. A set bit indicates that the associated pin is asserted.

This structure indicates the status of all the appropriate status bits at the time of the most recent error condition during a `write(2)` call, or the status of the bits at the most recent `BPPIOC_TESTIO` `ioctl(2)` call.

`pin_status` indicates possible error conditions under `ECPP_CENTRONICS` or `ECPP_COMPAT_MODE`. Under these modes, the state of the status pins will indicate the state of the device. For instance, many Centronics printers lower the `nErr` signal when a paper jam occurs. The behavior of the status pins depends on the device. Additional status information may be retrieved through the backchannel.

The `timeout_occurred` value is set when a timeout occurs during `write(2)`. `bus_error` is not used in this interface.

The following `ioctls` are used to directly read and write the parallel port status and control signals. If the current mode of the device is `ECPP_ECP_MODE` or `ECPP_NIBBLE_MODE`, the driver negotiates the link into `ECPP_COMPAT_MODE`, gets or sets the registers and then returns the link to `ECPP_NIBBLE_MODE`. If the current mode is `ECPP_CENTRONICS` or `ECPP_COMPAT_MODE`, these `ioctls` will get/set the register values in the current mode.

ECPPIOC_GETREGS

Read register values. The argument is a pointer to a `struct ecpp_regs`. See below for a description of this structure.

ecpp(7D)

	ECPPIOC_SETREGS	Set <code>ecpp</code> register values. The argument is a pointer to a <code>struct ecpp_regs</code> . See below for a description of this structure. If a parameter is out of range, <code>EINVAL</code> is returned.
		The Port Register Structure is defined in <code><sys/ecppio.h></code> .
		<pre>struct ecpp_regs { uchar dsr; /* status reg */ u_char dcr; /* control reg */ };</pre>
		The status register is read-only. The <code>ECPPIOC_SETREGS</code> ioctl has no affect on this register. Valid bit values for <code>dsr</code> are: <code>ECPP_nERR</code> , <code>ECPP_SLCT</code> , <code>ECPP_PE</code> , <code>ECPP_nACK</code> , <code>ECPP_nBUSY</code> . All other bits are reserved and always return 1.
		The control register is read/write. Valid bit values for <code>dcr</code> are: <code>ECPP_STB</code> , <code>ECPP_AFX</code> , <code>ECPP_nINIT</code> , <code>ECPP_SLCTIN</code> . All other bits are reserved. Reading reserved bits always return 1. An attempt to write 0s into these bits results in <code>EINVAL</code> .
DEVICE SPECIAL FILES	<code>/dev/lpN</code>	Solaris x86 only. (Backwards compatibility with former <code>lp(7D)</code> devices.)
	<code>/dev/printers/N</code>	1284 compliant parallel port device special files appears in both namespaces.
FILES	<code>kernel/drv/ecpp</code>	32-bit ELF kernel module
	<code>kernel/drv/sparcv9/ecpp</code>	64-bit SPARC ELF kernel module
	<code>kernel/drv/amd64/ecpp</code>	64-bit x86 ELF kernel module
	<code>kernel/drv/ecpp.conf</code>	driver configuration file
	<code>kernel/drv/sparcv9/ecpp.conf</code>	driver configuration file for 64-bit SPARC
	<code>kernel/drv/amd64/ecpp.conf</code>	driver configuration file for 64-bit x86
ERRORS	<code>EBADF</code>	The device is opened for write-only access and a read is attempted, or the device is opened for read-only access and a write is attempted.
	<code>EBUSY</code>	The device has been opened and another open is attempted. An attempt has been made to unload the driver while one of the units is open.
	<code>EINVAL</code>	A <code>ECPPIOC_SETPARMS</code> ioctl() is attempted with an out-of-range value in the <code>ecpp_transfer_parms</code> structure. A <code>ECPPIOC_SETREGS</code> ioctl() is attempted with an invalid value in the <code>ecpp_regs</code> structure. An ioctl() is attempted with an invalid value in the command argument. An invalid command argument is received during <code>modload(1M)</code> or <code>modunload(1M)</code> .

ecpp(7D)

EIO The driver encountered a bus error when attempting an access. A read or write did not complete properly, due to a peripheral error or a transfer timeout.

ENXIO The driver has received an open request for a unit for which the attach failed. The driver has received a write request for a unit which has an active peripheral error.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PCI-based systems
	ISA-based systems (x86)
Availability	SUNWpd (Sparc)
	SUNWpsdcr (x86)
Interface stability	Evolving

SEE ALSO `modload(1M)`, `modunload(1M)`, `update_drv(1M)`, `ioctl(2)`, `open(2)`, `read(2)`, `write(2)`, `attributes(5)`, `bpp(7D)`, `usbprn(7D)`, `prnio(7I)`, `streamio(7I)`

IEEE Std 1284-1994

<http://www.sun.com/io>

DIAGNOSTICS Parallel port controller not supported
Driver does not support parallel port controller on the given host. Attach failed.

ehci(7D)

NAME	ehci – Enhanced host controller driver
SYNOPSIS	usb@unit-address
DESCRIPTION	<p>The ehci driver is a USBA (Solaris USB Architecture) compliant nexus driver that supports the Enhanced Host Controller Interface Specification 2.0, an industry standard developed by Intel.</p> <p>The ehci supports bulk, interrupt, control and iso chronous transfers (on USB1.x devices behind a USB2.0 hub).</p>
FILES	<p>/kernel/drv/ehci 32-bit ELF kernel module x86 platform</p> <p>/kernel/drv/sparcv9/ehci 64-bit SPARC ELF kernel module</p> <p>/kernel/drv/amd64/ehci 64-bit x86 ELF kernel module</p> <p>/kernel/drv/ehci.conf Driver configuration file</p>

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86, PCI-based systems
Availability	SUNWusb

SEE ALSO add_drv(1M), prtconf(1M), update_drv(1M), attributes(5), hubd(7D), uhci(7D), ohci(7D), usba(7D)

Writing Device Drivers

Universal Serial Bus Specification 2.0

Enhanced Host Controller Interface Specification 1.0

System Administration Guide: Basic Administration

<http://www.sun.com/io>

DIAGNOSTICS All host controller errors are passed to the client drivers. Root errors are documented in hubd(7D).

In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

WARNING: <device path> <ehci><instance number>: Error message...

Due to recently discovered incompatibilities with this USB controller, USB2.x transfer support has been disabled. This device will continue to function as a USB1.x controller. If you are interested in enabling USB2.x support please refer to the ehci(7D) man page. Please refer to www.sun.com/io for Solaris Ready products and to www.sun.com/bigadmin/hcl for additional compatible USB products.

To enable USB2.x support, a new and more specific driver alias (refer to `add_drv(1M)` and `update_drv(1M)`) must be specified for ehci. By default, the ehci alias is 'pciclass,0c0320'. The compatible names in the `prtconf(1M)` output provides additional aliases. For example:

```
# prtconf -vp | grep pciclass,0c0320
      compatible: 'pci1106,3104.1106.3104.2063' + 'pci1106,3104.1106.3104' + 'pci1106,3104'
'pci1106,3104.2063' + 'pci1106,3104' + 'pciclass,0c0320' + 'pciclass,0c03'
....
```

A more specific alias is 'pci1106,3104'. Perform the following step to add this alias, then reboot the system:

```
# update_drv -a -i "pci1106,3104" ehci

# devfsadm -i ehci

or reboot
```

Connecting a high speed isochronous device to a high speed port is not supported. Devices requiring high speed isochronous USB service, such as high speed audio/video devices, are not supported.

Unrecoverable USB hardware error.

There was an unrecoverable USB hardware error reported by the ehci controller. Reboot the system. If this problem persists, contact your system vendor.

No SOF interrupts.

No SOF interrupts have been received. This USB EHCI controller is unusable.

The USB hardware is not generating Start Of Frame interrupts. Reboot the system. If this problem persists, contact your system vendor.

Please hotplug the 2.0 hub that this device is connected to.

A device error was detected. Remove and reinsert the external USB2.0 hub.

ehci_init_ctlr: Revision XX is not supported.

High speed USB devices < revision .95 are not supported.

ehci_init_ctlr: Unable to take control from BIOS.

High speed USB devices are not supported. In this case, Solaris was unable to take control of the USB hardware from the system's BIOS. Contact your system vendor or your system administrator for possible changes in BIOS settings. (x86 only.)

e1xl(7D)

NAME	e1xl – 3Com Ethernet device driver
SYNOPSIS	/kernel/drv/e1xl
DESCRIPTION	<p>The e1xl driver currently supports the following network cards: EtherLink XL (3C900-TPO, 3C900-COMBO, 3C900B-TPO, 3C900B-COMBO, and 3C900B-TPC), EtherLink XL 10/100 (3C905-TX Fast, 3C905-T4 Fast, 3C905B-TX Fast, 3C905B-T4 Fast, and 3C905C-TX-M Fast), and EtherLink Server 10/100 (3C980-TX Fast and 3C980C-TXM).</p> <p>The e1xl Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, dlpi(7P). Multiple EtherLink XL controllers installed within the system are supported by the driver. The e1xl driver provides basic support for the EtherLink hardware. Functions include chip initialization, frame transmit and receive, multicast and promiscuous mode support, and error recovery and reporting.</p> <p>The cloning, character-special device /dev/e1xl is used to access all EtherLink devices installed within the system.</p> <p>The e1xl driver is dependent on /kernel/misc/gld, a loadable kernel module that provides the e1xl driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on the primitives supported by the driver.</p> <p>The values returned by the driver in the DL_INFO_ACK primitive in response to the DL_INFO_REQ from the user are as follows:</p> <ul style="list-style-type: none">■ Maximum SDU is 1500 (ETHERMTU).■ Minimum SDU is 0. The driver will pad to the mandatory 60-octet minimum packet size.■ The dlsap address length is 8.■ MAC type is DL_ETHER.■ The sap length value is -2, meaning the physical address component is followed immediately by a 2-byte sap component within the DLSAP address.■ The broadcast address value is Ethernet/IEEE broadcast address (FF:FF:FF:FF:FF:FF).
PRECONFIGURATION	
Supported Settings	<ul style="list-style-type: none">■ Media Type: Auto Select
Known Problems and Limitations	<ul style="list-style-type: none">■ 3C905B cards in a Compaq ProLiant 6500 can fail to generate interrupts. There is no known workaround for this problem. However, because some slots appear to be more prone to the problem than others, try correcting the problem by moving the card to another PCI slot. If that fails, try rebooting the machine a number of times to free the card from the wedged state.■ Early versions of the 3Com 3C905C-TX-M adapter firmware do not support PXE network boot on Solaris systems. If you are using a version earlier than 4.11, upgrade the firmware. The PXE version is indicated by the Managed Boot Agent

version number. This number is not normally displayed during boot, but is shown on the PXE configuration screen.

CONFIGURATION	<p>The <code>/kernel/drv/elxl.conf</code> file supports the following tunable properties that you can set in the <code>elxl.conf</code> file:</p> <p><code>tx_start_thresh</code> Minimum number of bytes in transmit FIFO (if lower than the packet size), before NIC starts the frame transfer. The higher the value, the less chance for a transmission underrun error that triggers a transmission retry. Default value is 250, maximum value is 1514.</p> <p><code>recv-descriptors</code> The number of frames the device may receive from the network without system attention before the device begins dropping the frames. Valid range is 1 to 256. Default is 24.</p> <p><code>xmit-descriptors</code> The number of outgoing frames that the system can queue on the driver queue before blocking a send request. Valid range is 1 to 512. Default is 128.</p> <p><code>min-recv-data-buffers</code> Minimal number of allocated receive data buffers. Must be equal or greater than effective number of <code>recv-descriptors</code>. Default (max value) is 128.</p> <p><code>max-recv-data-buffers</code> Maximum number of allocated receive data buffers. Must be equal or greater than <code>min-recv-data-buffers</code> number.</p> <p><code>inter-frame-space</code> Amount of time in addition to the standard inter-frame time used by <i>IEEE 802.3</i> deference rule. Expressed in 32 multiples of bit time. Used only in half-duplex context. This small IFS is to make the defer-after-transmit time slightly longer than the defer-after-receive time. This greatly reduces the high collision rate seen with heavy TCP traffic (almost exactly one collision per ACK), giving slightly improved overall throughput in half-duplex mode. Default IFS is 32. This property is ignored in full-duplex mode.</p> <p><code>full-duplex</code> When set, it forces full-duplex mode for NIC capable of working in full-duplex mode. If this property is set, the auto-negotiation feature (if supported) of the network controller is disabled. For full duplex operation use <code>full-duplex=1</code>. For half duplex use <code>full-duplex=0</code>. Half-duplex operation provides better results on older 10-Mbit networks.</p> <p><code>speed</code> Link speed in Mbps. Valid values are 10 and 100. If this property is set, the auto-negotiation feature of the network controller (if supported) is disabled.</p>				
FILES	<table border="0"> <tr> <td style="padding-right: 20px;"><code>/dev/elxl</code></td> <td>Special character device</td> </tr> <tr> <td><code>/kernel/drv/elxl.conf</code></td> <td>Configuration file for <code>elxl</code> driver</td> </tr> </table>	<code>/dev/elxl</code>	Special character device	<code>/kernel/drv/elxl.conf</code>	Configuration file for <code>elxl</code> driver
<code>/dev/elxl</code>	Special character device				
<code>/kernel/drv/elxl.conf</code>	Configuration file for <code>elxl</code> driver				

elxl(7D)

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO [attributes\(5\)](#), [gld\(7D\)](#), [streamio\(7I\)](#), [dlpi\(7P\)](#).

NAME	eri – eri Fast-Ethernet device driver
SYNOPSIS	<code>/dev/eri</code>
DESCRIPTION	<p>The <code>eri</code> Fast Ethernet driver is a multi-threaded, loadable, clonable, STREAMS—based hardware driver supporting the connectionless Data Link Provider Interface <code>dlpi(7P)</code> over an <code>eri</code> Fast-Ethernet controller. Multiple <code>eri</code> devices installed within the system are supported by the driver.</p> <p>The <code>eri</code> driver provides basic support for the <code>eri</code> hardware and handles the <code>eri</code> device. Functions include chip initialization, frame transit and receive, multicast and promiscuous support, and error recovery and reporting.</p> <p>The <code>eri</code> device provides 100Base-TX networking interfaces using the SUN <code>RIO</code> ASIC and an internal transceiver. The <code>RIO</code> ASIC provides the PCI interface and MAC functions. The physical layer functions are provided by the internal transceiver which connects to a RJ-45 connector.</p> <p>The 100Base-TX standard specifies an auto-negotiation protocol to automatically select the mode and speed of operation. The internal transceiver is capable of performing auto-negotiation using the remote-end of the link (link partner) and receives the capabilities of the remote end. It selects the highest common denominator mode of operation based on the priorities. It also supports a forced-mode of operation under which the driver selects the mode of operation.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device <code>/dev/eri</code> is used to access all <code>eri</code> controllers installed within the system.</p>
eri and DLPI	<p>The <code>eri</code> driver is a “style 2” Data Link Service provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code> type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <code><sys/dlpi.h></code>. Refer to <code>dlpi(7P)</code> for more information.</p> <p>An explicit <code>DL_ATTACH_REQ</code> message by the user is required to associate the opened stream with a particular device (<code>ppa</code>). The <code>ppa</code> ID is interpreted as an unsigned integer data type and indicates the corresponding device instance (unit) number. An error (<code>DL_ERROR_ACK</code>) is returned by the driver if the <code>ppa</code> field value does not correspond to a valid device instance number for this system. The device is initialized on first attach and de-initialized (stopped) at last detach.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> from the user are as follows:</p> <ul style="list-style-type: none"> ■ The maximum SDU is 1500 (<code>ETHERMTU</code> - defined in <code><sys/ethernet.h></code>). ■ The minimum SDU is 0. ■ The <code>dlsap</code> address length is 8. ■ The MAC type is <code>DL_ETHER</code>. ■ The <code>sap</code> length values is -2, meaning the physical address component is followed immediately by a 2 byte <code>sap</code> component within the DLSAP address. ■ The service mode is <code>DL_CLDLS</code>.

eri(7D)

- Optional quality of service (QOS) is not currently supported so QOS fields are 0.
- The provider style is `DL_STYLE`.
- The version is `DL_VERSION_2`.
- The broadcast address value is Ethernet/IEEE broadcast address (`0xFFFFFFFF`).

Once in the `DL_ATTACHED` state, the user must send a `DL_BIND_REQ` to associate a particular SAP (Service Access Pointer) with the stream. The `eri` driver interprets the `sap` field within the `DL_BIND_REQ` as an Ethernet "type," therefore valid values for the `sap` field are in the `[0-0xFFFF]` range. Only one Ethernet type can be bound to the stream at any time.

If the user selects a `sap` with a value of 0, the receiver will be in IEEE 802.3 mode. All frames received from the media having a Ethernet type field in the range `[0-1500]` are assumed to be 802.3 frames and are routed up all open Streams which are bound to `sap` value 0. If more than one Stream is in 802.3 mode, the frame will be duplicated and routed up multiple Streams as `DL_UNITDATA_IND` messages.

In transmission, the driver checks the `sap` field of the `DL_BIND_REQ` to determine if the value is 0 or if the Ethernet type field is in the range `[0-1500]`. If either is true, the driver computes the length of the message, not including initial `M_PROTO` mblk (message block), of all subsequent `DL_UNITDATA_REQ` messages, and transmits 802.3 frames that have this value in the MAC frame header length field.

The `eri` driver's DLSAP address format consists of the 6 byte physical (Ethernet) address component followed immediately by the 2 byte `sap` (type) component, producing an 8 byte DLSAP address. Applications should *not* hardcode to this particular implementation-specific DLSAP address format but use information returned in the `DL_INFO_ACK` primitive to compose and decompose DLSAP addresses. The `sap` length, full DLSAP length, and `sap`/physical ordering are included within the `DL_INFO_ACK`. The physical address length can be computed by subtracting the `sap` length from the full DLSAP address length or by issuing the `DL_PHYS_ADDR_REQ` to obtain the current physical address associated with the stream.

Once in the `DL_BOUND` state, the user may transmit frames on the Ethernet by sending `DL_UNITDATA_REQ` messages to the `eri` driver. The `eri` driver will route received Ethernet frames up all open and bound streams having a `sap` which matches the Ethernet type as `DL_UNITDATA_IND` messages. Received Ethernet frames are duplicated and routed up multiple open streams if necessary. The DLSAP address contained within the `DL_UNITDATA_REQ` and `DL_UNITDATA_IND` messages consists of both the `sap` (type) and physical (Ethernet) components.

eri Primitives

In addition to the mandatory connectionless DLPI message set, the driver also supports the following primitives:

The `DL_ENABMULTI_REQ` and `DL_DISABMULTI_REQ` primitives enable/disable reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. These primitives are accepted by the driver in any state following `DL_ATTACHED`.

The `DL_PROMISCON_REQ` and `DL_PROMISCOFF_REQ` primitives with the `DL_PROMISC_PHYS` flag set in the `d1_level` field enables/disables reception of all promiscuous mode frames on the media, including frames generated by the local host. When used with the `DL_PROMISC_SAP` flag set, this enables/disables reception of all `sap` (Ethernet type) values. When used with the `DL_PROMISC_MULTI` flag set, this enables/disables reception of all multicast group addresses. The effect of each is always on a per-stream basis and independent of the other `sap` and physical level configurations on this stream or other streams.

The `DL_PHYS_ADDR_REQ` primitive returns the 6 octet Ethernet address currently associated (attached) to the stream in the `DL_PHYS_ADDR_ACK` primitive. This primitive is valid only in states following a successful `DL_ATTACH_REQ`.

The `DL_SET_PHYS_ADDR_REQ` primitive changes the 6 octet Ethernet address currently associated (attached) to this stream. The credentials of the process which originally opened this stream must be superuser, or `EPERM` is returned in the `DL_ERROR_ACK`. This primitive is destructive because it affects all current and future streams attached to this device. An `M_ERROR` is sent up all other streams attached to this device when this primitive is successful on this stream. Once changed, all streams subsequently opened and attached to this device will obtain this new physical address. Once changed, the physical address will remain until this primitive is used to change the physical address again or the system is rebooted, whichever comes first.

eri DRIVER

By default, the `eri` driver performs auto-negotiation to select the mode and speed of the link, which can be in one of the following modes, as described in the 100Base-TX standard:

- 100 Mbps, full-duplex
- 100 Mbps, half-duplex
- 10 Mbps, full-duplex
- 10 Mbps, half-duplex

The auto-negotiation protocol automatically selects:

- Operation mode (half-duplex or full-duplex)
- Speed (100 Mbps or 10 Mbps)

The auto-negotiation protocol does the following:

- Gets all modes of operation supported by the link partner
- Advertises its capabilities to the Link Partner
- Selects the highest common denominator mode of operation based on the priorities

The internal transceiver is capable of all of the operating speeds and modes listed above. By default, auto-negotiation is used to select the speed and the mode of the link and the common mode of operation with the link partner.

For users who want to select the speed and mode of the link, the `eri` device supports programmable IPG (Inter-Packet Gap) parameters `ipg1` and `ipg2`. Sometimes, the user may want to alter these values depending on whether the driver supports 10 Mbps or 100 Mbps and accordingly, IPG will be set to 9.6 or 0.96 microseconds.

eri(7D)

eri Parameter List

The `eri` driver provides for setting and getting various parameters for the `eri` device. The parameter list includes current transceiver status, current link status, inter-packet gap, local transceiver capabilities and link partner capabilities.

The local transceiver has two set of capabilities: one set reflects hardware capabilities, which are read-only (RO) parameters. The second set reflects the values chosen by the user and is used in speed selection and possess read/write (RW) capability. At boot time, these two sets of capabilities will be the same. Because the current default value of these parameters can only be read and not modified, the link partner capabilities are also read only.

FILES	<code>/dev/eri</code>	<code>eri</code> special character device.
	<code>/kernel/drv/eri.conf</code>	System wide default device driver properties
	<code>/kernel/drv/sparcv9/eri</code>	64 bit device driver

SEE ALSO `ndd(1M)`, `netstat(1M)`, `driver.conf(4)`, `hme(7D)`, `qfe(7D)`, `dlpi(7P)`

NAME	esp – ESP SCSI Host Bus Adapter Driver
SYNOPSIS	<code>esp@sbus-slot, 80000</code>
DESCRIPTION	<p>The <code>esp</code> Host Bus Adapter driver is a SCSI compliant nexus driver that supports the Emulex family of esp SCSI chips (<code>esp100</code>, <code>esp100A</code>, <code>esp236</code>, <code>fas101</code>, <code>fas236</code>).</p> <p>The <code>esp</code> driver supports the standard functions provided by the SCSI interface. The driver supports tagged and untagged queuing, fast SCSI (on FAS esp's only), almost unlimited transfer size (using a moving DVMA window approach), and auto request sense; but it does not support linked commands.</p>
CONFIGURATION	<p>The <code>esp</code> driver can be configured by defining properties in <code>esp.conf</code> which override the global SCSI settings. Supported properties are: <code>scsi-options</code>, <code>target<n>-scsi-options</code>, <code>scsi-reset-delay</code>, <code>scsi-watchdog-tick</code>, <code>scsi-tag-age-limit</code>, <code>scsi-initiator-id</code>.</p> <p><code>target<n>-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target<n></code>. <code><n></code> can vary from 0 to 7.</p> <p>Refer to <code>scsi_hba_attach(9F)</code> for details.</p>
EXAMPLES	<p>EXAMPLE 1 A sample of <code>esp</code> configuration file.</p> <p>Create a file <code>/kernel/drv/esp.conf</code> and add this line:</p> <pre>scsi-options=0x78;</pre> <p>This will disable tagged queuing, fast SCSI, and Wide mode for all <code>esp</code> instances. To disable an option for one specific <code>esp</code> (refer to <code>driver.conf(4)</code>):</p> <pre>name="esp" parent="/iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000" reg=0xf,0x800000,0x40 target1-scsi-options=0x58 scsi-options=0x178 scsi-initiator-id=6;</pre> <p>Note that the default initiator ID in OBP is 7 and that the change to ID 6 will occur at attach time. It may be preferable to change the initiator ID in OBP.</p> <p>The above would set <code>scsi-options</code> for target 1 to 0x58 and for all other targets on this SCSI bus to 0x178. The physical pathname of the parent can be determined using the <code>/devices</code> tree or following the link of the logical device name:</p> <pre>example# ls -l /dev/rdisk/c0t3d0s0 lrwxrwxrwx 1 root root 88 Aug 22 13:29 /dev/rdisk/c0t3d0s0 -> ../../../../devices/iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/ esp@f,800000/sd@3,0:a,raw</pre> <p>The register property values can be determined from <code>prtconf(1M)</code> output (<code>-v</code> option):</p> <pre>esp, instance #0 Register Specifications: Bus Type=0xf, Address=0x800000, Size=40</pre>

esp(7D)

EXAMPLE 1 A sample of esp configuration file. (Continued)

To set scsi-options more specifically per target:

```
target1-scsi-options=0x78;  
device-type-scsi-options-list =  
    "SEAGATE ST32550W", "seagate-scsi-options" ;  
seagate-scsi-options = 0x58;  
scsi-options=0x3f8;
```

The above would set scsi-options for target 1 to 0x78 and for all other targets on this SCSI bus to 0x378 except for one specific disk type which will have scsi-options set to 0x58.

scsi-options specified per target ID has the highest precedence, followed by scsi-options per device type. To get the inquiry string run probe-scsi or probe-scsi-all command at the ok prompt before booting the system.

Global, for example. for all esp instances, scsi-options per bus has the lowest precedence.

The system needs to be rebooted before the specified scsi-options take effect.

FILES /kernel/drv/esp ELF Kernel Module
/kernel/drv/esp.conf Configuration file

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SBus-based systems with esp-based
	SCSI port and SSHA, SBE/S, FSBE/S, and DSBE/S SBus SCSI Host Adapter options

SEE ALSO prtconf(1M), driver.conf(4), attributes(5), [fas\(7D\)](#), scsi_abort(9F), scsi_hba_attach(9F), scsi_ifgetcap(9F), scsi_reset(9F), scsi_sync_pkt(9F), scsi_transport(9F), scsi_device(9S), scsi_extended_sense(9S), scsi_inquiry(9S), scsi_pkt(9S)

Writing Device Drivers

OpenBoot Command Reference

ANSI Small Computer System Interface-2 (SCSI-2)

ESP Technical Manuals, QLogic Corp.

DIAGNOSTICS The messages described below are some that may appear on the system console, as well as being logged.

The first four messages may be displayed while the esp driver is trying to attach; these messages mean that the esp driver was unable to attach. All of these messages are preceded by "esp%d", where "%d" is the instance number of the esp controller.

Device in slave-only slot

The SBus device has been placed in a slave-only slot and will not be accessible; move to non-slave-only SBus slot.

Device is using a hilevel intr

The device was configured with an interrupt level that cannot be used with this esp driver. Check the SBus device.

Unable to map registers

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot find dma controller

Driver was unable to locate a dma controller. This is an auto-configuration error.

Disabled TQ since disconnects are disabled

Tagged queuing was disabled because disconnects were disabled in scsi-options.

Bad clock frequency- setting 20mhz, asynchronous mode

Check for bad hardware.

Sync pkt failed

Syncing a SCSI packet failed. Refer to scsi_sync_pkt(9F).

Slot %x: All tags in use!!!

The driver could not allocate another tag number. The target devices do not properly support tagged queuing.

Target %d.%d cannot alloc tag queue\n

The driver could not allocate space for tag queue.

Gross error in esp status (%x)

The driver experienced severe SCSI bus problems. Check cables and terminator.

Spurious interrupt

The driver received an interrupt while the hardware was not interrupting.

Lost state in phasemanage

The driver is confused about the state of the SCSI bus.

Unrecoverable DMA error during selection

The DMA controller experienced host SBus problems. Check for bad hardware.

Bad sequence step (0x%x) in selection

The esp hardware reported a bad sequence step. Check for bad hardware.

Undetermined selection failure

The selection of a target failed unexpectedly. Check for bad hardware.

>2 reselection IDs on the bus

Two targets selected simultaneously, which is illegal. Check for bad hardware.

esp(7D)

Reconnect: unexpected bus free
A reconnect by a target failed. Check for bad hardware.

Timeout on receiving tag msg
Suspect target f/w failure in tagged queue handling.

Parity error in tag msg
A parity error was detected in a tag message. Suspect SCSI bus problems.

Botched tag
The target supplied bad tag messages. Suspect target f/w failure in tagged queue handling.

Parity error in reconnect msg's
The reconnect failed because of parity errors.

Target <n> didn't disconnect after sending <message>
The target unexpectedly did not disconnect after sending <message>.

No support for multiple segs
The esp driver can only transfer contiguous data.

No dma window?
Moving the DVMA window failed unexpectedly.

No dma window on <type> operation
Moving the DVMA window failed unexpectedly.

Cannot set new dma window
Moving the DVMA window failed unexpectedly.

Unable to set new window at <address> for <type> operation
Moving the DVMA window failed unexpectedly.

Illegal dma boundary? %x
An attempt was made to cross a boundary that the driver could not handle.

Unwanted data out/in for Target <n>
The target went into an unexpected phase.

Spurious <name> phase from target <n>
The target went into an unexpected phase.

SCSI bus DATA IN phase parity error
The driver detected parity errors on the SCSI bus.

SCSI bus MESSAGE IN phase parity error
The driver detected parity errors on the SCSI bus.

SCSI bus STATUS phase parity error
The driver detected parity errors on the SCSI bus.

Premature end of extended message
An extended SCSI bus message did not complete. Suspect a target f/w problem.

Premature end of input message
A multibyte input message was truncated. Suspect a target f/w problem.

Input message botch

The driver is confused about messages coming from the target.

Extended message <n> is too long

The extended message sent by the target is longer than expected.

<name> message <n> from Target <m> garbled

Target <m> sent message <name> of value <n> which the driver did not understand.

Target <n> rejects our message <name>

Target <n> rejected a message sent by the driver.

Rejecting message <name> from Target <n>

The driver rejected a message received from target <n>

Cmd dma error

The driver was unable to send out command bytes.

Target <n> refused message resend

The target did not accept a message resend.

Two-byte message <name> <value> rejected

The driver does not accept this two-byte message.

Unexpected selection attempt

An attempt was made to select this host adapter by another initiator.

Polled cmd failed (target busy)

A polled command failed because the target did not complete outstanding commands within a reasonable time.

Polled cmd failed

A polled command failed because of timeouts or bus errors.

Disconnected command timeout for Target <id>.<lun>

A timeout occurred while target/lun was disconnected. This is usually a target f/w problem. For tagged queuing targets, <n> commands were outstanding when the timeout was detected.

Disconnected tagged cmds (<n>) timeout for Target <id>.<lun>

A timeout occurred while target/lun was disconnected. This is usually a target f/w problem. For tagged queuing targets, <n> commands were outstanding when the timeout was detected.

Connected command timeout for Target <id>.<lun>

This is usually a SCSI bus problem. Check cables and termination.

Target <id>.<lun> reverting to async. mode

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id>.<lun> reducing sync. transfer rate

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

esp(7D)

Reverting to slow SCSI cable mode

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Reset SCSI bus failed

An attempt to reset the SCSI bus failed.

External SCSI bus reset

Another initiator reset the SCSI bus.

WARNINGS

The esp hardware does not support Wide SCSI mode. Only FAS-type esp's support fast SCSI (10 MB/sec).

NOTES

The esp driver exports properties indicating per target the negotiated transfer speed (target<n>-sync-speed) and whether tagged queuing has been enabled (target<n>-TQ). The sync-speed property value is the data transfer rate in KB/sec. The target-TQ property has no value. The existence of the property indicates that tagged queuing has been enabled. Refer to prtconf(1M) (verbose option) for viewing the esp properties.

```
dma, instance #3
  Register Specifications:
    Bus Type=0x2, Address=0x81000, Size=10
esp, instance #3
  Driver software properties:
    name <target3-TQ> length <0> - <no
value>.
    name <target3-sync-speed> length <4>
    value <0x00002710>.
    name <scsi-options> length <4>
    value <0x000003f8>.
    name <scsi-watchdog-tick> length <4>
    value <0x0000000a>.
    name <scsi-tag-age-limit> length <4>
    value <0x00000008>.
    name <scsi-reset-delay> length <4>
    value <0x00000bb8>.
```

NAME	fas – FAS SCSI Host Bus Adapter Driver
SYNOPSIS	<code>fas@sbus-slot, 0x8800000</code>
DESCRIPTION	<p>The <code>fas</code> Host Bus Adapter driver is a SCSI compliant nexus driver that supports the Qlogic FAS366 SCSI chip.</p> <p>The <code>fas</code> driver supports the standard functions provided by the SCSI interface. The driver supports tagged and untagged queuing, wide and fast SCSI, almost unlimited transfer size (using a moving DVMA window approach), and auto request sense; but it does not support linked commands.</p>
Driver Configuration	<p>The <code>fas</code> driver can be configured by defining properties in <code>fas.conf</code> which override the global SCSI settings. Supported properties are: <code>scsi-options</code>, <code>target<n>-scsi-options</code>, <code>scsi-reset-delay</code>, <code>scsi-watchdog-tick</code>, <code>scsi-tag-age-limit</code>, <code>scsi-initiator-id</code>.</p> <p><code>target<n>-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target<n></code>. <code><n></code> can vary from decimal 0 to 15. The supported <code>scsi-options</code> are: <code>SCSI_OPTIONS_DR</code>, <code>SCSI_OPTIONS_SYNC</code>, <code>SCSI_OPTIONS_TAG</code>, <code>SCSI_OPTIONS_FAST</code>, and <code>SCSI_OPTIONS_WIDE</code>.</p> <p>After periodic interval <code>scsi-watchdog-tick</code>, the <code>fas</code> driver searches all current and disconnected commands for timeouts.</p> <p><code>scsi-tag-age-limit</code> is the number of times that the <code>fas</code> driver attempts to allocate a particular tag ID that is currently in use after going through all tag IDs in a circular fashion. After finding the same tag ID in use <code>scsi-tag-age-limit</code> times, no more commands will be submitted to this target until all outstanding commands complete or timeout.</p> <p>Refer to <code>scsi_hba_attach(9F)</code> for details.</p>
EXAMPLES	<p>EXAMPLE 1 A sample of <code>fas</code> configuration file</p> <p>Create a file called <code>/kernel/drv/fas.conf</code> and add this line:</p> <pre>scsi-options=0x78;</pre> <p>This disables tagged queuing, Fast SCSI, and Wide mode for all <code>fas</code> instances. The following example disables an option for one specific <code>fas</code> (refer to <code>driver.conf(4)</code> for more details):</p> <pre>name="fas" parent="/iommu@f,e0000000/sbus@f,e0001000" reg=3,0x8800000,0x10,3,0x8810000,0x40 target1-scsi-options=0x58 scsi-options=0x178 scsi-initiator-id=6;</pre> <p>Note that the default initiator ID in OBP is 7 and that the change to ID 6 will occur at attach time. It may be preferable to change the initiator ID in OBP.</p> <p>The example above sets <code>scsi-options</code> for target 1 to 0x58 and all other targets on this SCSI bus to 0x178.</p>

fas(7D)

The physical pathname of the parent can be determined using the `/devices` tree or following the link of the logical device name:

```
# ls -l /dev/rdisk/c1t3d0s0
lrwxrwxrwx 1 root  other  78 Aug 28 16:05 /dev/rdisk/c1t3d0s0 ->
. . / . . /devices/iommu@f,e0000000\
      sbus@f,e0001000/SUNW,fas@3,8800000/sd@3,0:a,raw
```

Determine the register property values using the output from `prtconf(1M)` (with the `-v` option):

```
SUNW,fas, instance #0
```

```
Register Specifications:
```

```
Bus Type=0x3, Address=0x8800000, Size=10
```

```
Bus Type=0x3, Address=0x8810000, Size=40
```

`scsi-options` can also be specified per device type using the device inquiry string. All the devices with the same inquiry string will have the same `scsi-options` set. This can be used to disable some `scsi-options` on all the devices of the same type.

```
device-type-scsi-options-list=
```

```
"TOSHIBA XM5701TASUN12XCD", "cd-scsi-options";
```

```
cd-scsi-options = 0x0;
```

The above entry in `/kernel/drv/fas.conf` sets the `scsi-options` for all devices with inquiry string `TOSHIBA XM5701TASUN12XCD` to `cd-scsi-options`. To get the inquiry string, run the `probe-scsi` or `probe-scsi-all` command at the `ok` prompt before booting the system.

To set `scsi-options` more specifically per target:

```
target1-scsi-options=0x78;
```

```
device-type-scsi-options-list =
```

```
"SEAGATE ST32550W", "seagate-scsi-options" ;
```

```
seagate-scsi-options = 0x58;
```

```
scsi-options=0x3f8;
```

The above sets `scsi-options` for target 1 to `0x78` and for all other targets on this SCSI bus to `0x3f8` except for one specific disk type which will have `scsi-options` set to `0x58`.

`scsi-options` specified per target ID have the highest precedence, followed by `scsi-options` per device type. Global `fas scsi-options` (effecting all instances) per bus have the lowest precedence.

The system needs to be rebooted before the specified `scsi-options` take effect.

Driver Capabilities

The target driver needs to set capabilities in the `fas` driver in order to enable some driver features. The target driver can query and modify these capabilities:

`synchronous`, `tagged-qing`, `wide-xfer`, `auto-rqsense`, `qfull-retries`, `qfull-retry-interval`. All other capabilities can only be queried.

By default, `tagged-qing`, `auto-rqsense`, and `wide-xfer` capabilities are disabled, while `disconnect`, `synchronous`, and `untagged-qing` are enabled. These capabilities can only have binary values (0 or 1). The default value for `qfull-retries` is 10 and the default value for `qfull-retry-interval` is 100. The `qfull-retries` capability is a `uchar_t` (0 to 255) while `qfull-retry-interval` is a `ushort_t` (0 to 65535).

The target driver needs to enable `tagged-qing` and `wide-xfer` explicitly. The `untagged-qing` capability is always enabled and its value cannot be modified, because `fas` can queue commands even when `tagged-qing` is disabled.

Whenever there is a conflict between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only `whom != 0` is supported in the `scsi_ifsetcap(9F)` call.

Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

FILES /kernel/drv/fas ELF Kernel Module
/kernel/drv/fas.conf Optional configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Limited to Sparc SBus-based systems with FAS366-based SCSI port and SunSWIFT SBus SCSI Host Adapter/Fast Ethernet option.

SEE ALSO `prtconf(1M)`, `driver.conf(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

ANSI Small Computer System Interface-2 (SCSI-2)

QLogic Corporation, *FAS366 Technical Manuals*.

DIAGNOSTICS The messages described below are some that may appear on the system console, as well as being logged.

The first five messages may be displayed while the `fas` driver is trying to attach; these messages mean that the `fas` driver was unable to attach. All of these messages are preceded by "`fas%d`", where "`%d`" is the instance number of the `fas` controller.

Device in slave-only slot

The SBus device has been placed in a slave-only slot and will not be accessible; move to non-slave-only SBus slot.

fas(7D)

Device is using a hilevel intr
The device was configured with an interrupt level that cannot be used with this fas driver. Check the SBus device.

Cannot alloc dma handle
Driver was unable to allocate memory for a DMA controller.

Cannot alloc cmd area
Driver was unable to allocate memory for a command address.

Cannot create kmem_cache
Driver was unable to allocate memory for internal data structures.

Unable to map FAS366 registers
Driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot add intr
Driver could not add its interrupt service routine to the kernel.

Cannot map dma
Driver was unable to locate a DMA controller. This is an auto-configuration error.

Cannot bind cmdarea
Driver was unable to bind the DMA handle to an address.

Cannot create devctl minor node
Driver is unable to create a minor node for the controller.

Cannot attach
The driver was unable to attach; usually follows another warning that indicates why attach failed.

Disabled TQ since disconnects are disabled
Tagged queuing was disabled because disconnects were disabled in `scsi-options`.

Bad clock frequency
Check for bad hardware.

Sync of pkt (<address>) failed
Syncing a SCSI packet failed. Refer to `scsi_sync_pkt(9F)`.

All tags in use!
The driver could not allocate another tag number. The target devices do not properly support tagged queuing.

Gross error in FAS366 status
The driver experienced severe SCSI bus problems. Check cables and terminator.

Spurious interrupt
The driver received an interrupt while the hardware was not interrupting.

Lost state in phasemanage
The driver is confused about the state of the SCSI bus.

Unrecoverable DMA error during selection
The DMA controller experienced host SBus problems. Check for bad hardware.

Bad sequence step (<step number>) in selection
The FAS366 hardware reported a bad sequence step. Check for bad hardware.

Undetermined selection failure
The selection of a target failed unexpectedly. Check for bad hardware.

Target <n>: failed reselection (bad reselect bytes)
A reconnect failed, target sent incorrect number of message bytes. Check for bad hardware.

Target <n>: failed reselection (bad identify message)
A reconnect failed, target didn't send identify message or it got corrupted. Check for bad hardware.

Target <n>: failed reselection (not in msgin phase)
Incorrect SCSI bus phase after reconnection. Check for bad hardware.

Target <n>: failed reselection (unexpected bus free)
Incorrect SCSI bus phase after reconnection. Check for bad hardware.

Target <n>: failed reselection (timeout on receiving tag msg)
A reconnect failed; target failed to send tag bytes. Check for bad hardware.

Target <n>: failed reselection (botched tag)
A reconnect failed; target failed to send tag bytes. Check for bad hardware.

Target <n>: failed reselection (invalid tag)
A reconnect failed; target sent incorrect tag bytes. Check for bad hardware.

Target <n>: failed reselection (Parity error in reconnect msg's)
A reconnect failed; parity error detected. Check for bad hardware.

Target <n>: failed reselection (no command)
A reconnect failed; target accepted `abort` or `reset`, but still tries to reconnect. Check for bad hardware.

Unexpected bus free
Target disconnected from the bus without notice. Check for bad hardware.

Target <n> didn't disconnect after sending <message>
The target unexpectedly did not disconnect after sending *<message>* .

Bad sequence step (0x?) in selection
The sequence step register shows an improper value. The target might be misbehaving.

Illegal dma boundary?
An attempt was made to cross a boundary that the driver could not handle.

Unwanted data xfer direction for Target <n>
The target went into an unexpected phase.

fas(7D)

Unrecoverable DMA error on dma <send/receive>
There is a DMA error while sending/receiving data. The host DMA controller is experiencing some problems.

SCSI bus DATA IN phase parity error
The driver detected parity errors on the SCSI bus.

SCSI bus MESSAGE IN phase parity error
The driver detected parity errors on the SCSI bus.

SCSI bus STATUS phase parity error
The driver detected parity errors on the SCSI bus.

Premature end of extended message
An extended SCSI bus message did not complete. Suspect a target firmware problem.

Premature end of input message
A multibyte input message was truncated. Suspect a target firmware problem.

Input message botch
The driver is confused about messages coming from the target.

Extended message <n> is too long
The extended message sent by the target is longer than expected.

<name> message <n> from Target <m> garbled
Target <m> sent message <name> of value <n> which the driver did not understand.

Target <n> rejects our message <name>
Target <n> rejected a message sent by the driver.

Rejecting message <name> from Target <n>
The driver rejected a message received from target <n> .

Cmd transmission error
The driver was unable to send out command bytes.

Target <n> refused message resend
The target did not accept a message resend.

MESSAGE OUT phase parity error
The driver detected parity errors on the SCSI bus.

Two byte message <name> <value> rejected
The driver does not accept this two byte message.

Gross error in fas status <stat>
The fas chip has indicated a gross error like FIFO overflow.

Polled cmd failed (target busy)
A polled command failed because the target did not complete outstanding commands within a reasonable time.

Polled cmd failed
A polled command failed because of timeouts or bus errors.

Auto request sense failed

Driver is unable to get request sense from the target.

Disconnected command timeout for Target <id>.<lun>

A timeout occurred while target *id/lun* was disconnected. This is usually a target firmware problem. For tagged queuing targets, <n> commands were outstanding when the timeout was detected.

Disconnected tagged cmds (<n>) timeout for Target <id>.<lun>

A timeout occurred while target *id/lun* was disconnected. This is usually a target firmware problem. For tagged queuing targets, <n> commands were outstanding when the timeout was detected.

Connected command timeout for Target <id>.<lun>

This is usually a SCSI bus problem. Check cables and termination.

Target <id>.<lun> reverting to async. mode

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id>.<lun> reducing sync. transfer rate

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Reverting to slow SCSI cable mode

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id> reducing sync. transfer rate

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id> reverting to async. mode

A data transfer hang was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id> disabled wide SCSI mode

Due to problems on the SCSI bus, the driver goes into more conservative mode of operation to avoid further problems.

Reset SCSI bus failed

An attempt to reset the SCSI bus failed.

External SCSI bus reset

Another initiator reset the SCSI bus.

WARNINGS

The *fas* hardware (FAS366) supports both Wide and Fast SCSI mode, but *fast20* is not supported. The maximum SCSI bandwidth is 20 MB/sec. Initiator mode block sequence (IBS) is not supported.

NOTES

The *fas* driver exports properties indicating per target the negotiated transfer speed (*target<n>-sync-speed*), whether wide bus is supported (*target<n>-wide*), *scsi-options* for that particular target (*target<n>-scsi-options*), and whether tagged queuing has been enabled (*target<n>-TQ*). The *sync-speed* property value

fas(7D)

is the data transfer rate in KB/sec. The target<n>-TQ and the target<n>-wide property have value 1 to indicate that the corresponding capability is enabled, or 0 to indicate that the capability is disabled for that target. Refer to prtconf(1M) (verbose option) for viewing the fas properties.

```
SUNW,fas,instance #1
Driver software properties:
  name <target3-TQ> length <4>
    value <0x00000001>.
  name <target3-wide> length <4>
    value <0x00000000>.
  name <target3-sync-speed> length <4>
    value <0x00002710>.
  name <target3-scsi-options> length <4>
    value <0x000003f8>.
  name <target0-TQ> length <4>
    value <0x00000001>.
  name <pm_norm_pwr> length <4>
    value <0x00000001>.
  name <pm_timestamp> length <4>
    value <0x30040346>.
  name <scsi-options> length <4>
    value <0x000003f8>.
  name <scsi-watchdog-tick> length <4>
    value <0x0000000a>.
  name <scsi-tag-age-limit> length <4>
    value <0x00000002>.
  name <scsi-reset-delay> length <4>
    value <0x00000bb8>.
Register Specifications:
  Bus Type=0x3, Address=0x8800000, Size=10
  Bus Type=0x3, Address=0x8810000, Size=40
Interrupt Specifications:
  Interrupt Priority=0x35 (ipl 5)
```

NAME	fasttrap – DTrace user instruction tracing provider
DESCRIPTION	<p>The <code>fasttrap</code> driver is a DTrace dynamic tracing provider that performs dynamic instrumentation of arbitrary instructions in Solaris processes. The <code>fasttrap</code> driver implements the DTrace <code>fasttrap</code> and <code>pid</code> providers.</p> <p>The <code>fasttrap</code> driver is not a public interface and you access instrumentation offered by this provider through DTrace. Refer to the <i>Solaris Dynamic Tracing Guide</i> for a description of the public documented interfaces available for the DTrace facility and the probes offered by the <code>fasttrap</code> provider.</p>
SPARC ONLY	The <code>fasttrap</code> provider provides a DTrace probe that fires each time a user process executes an instruction. The <code>pid</code> provider allows for the dynamic creation of DTrace probes corresponding to instruction locations inside any user process specified using a process ID and an instruction address or symbol name. Together these providers permit DTrace users to perform instrumentation of Solaris user processes and to trace the interactions between processes and the operating system. See the chapter entitled "User Process Tracing" in the <i>Solaris Dynamic Tracing Guide</i> for information on how to use these providers to instrument processes.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdtrp
Interface Stability	Private

SEE ALSO `dtrace(1M)`, `attributes(5)`, `dtrace(7D)`
Solaris Dynamic Tracing Guide

fbio(7I)

NAME	fbio – frame buffer control operations
DESCRIPTION	<p>The frame buffers provided with this release support the same general interface that is defined by <code><sys/fbio.h></code>. Each responds to an <code>FBIOGTYPE ioctl(2)</code> request which returns information in a <code>fbtype</code> structure.</p> <p>Each device has an <code>FBTYPE</code> which is used by higher-level software to determine how to perform graphics functions. Each device is used by opening it, doing an <code>FBIOGTYPE ioctl()</code> to see which frame buffer type is present, and thereby selecting the appropriate device-management routines.</p> <p><code>FBIOGINFO</code> returns information specific to the GS accelerator.</p> <p><code>FBIOSVIDEO</code> and <code>FBIOGVIDEO</code> are general-purpose <code>ioctl()</code> requests for controlling possible video features of frame buffers. These <code>ioctl()</code> requests either set or return the value of a flags integer. At this point, only the <code>FBVIDEO_ON</code> option is available, controlled by <code>FBIOSVIDEO</code>. <code>FBIOGVIDEO</code> returns the current video state.</p> <p>The <code>FBIOSATTR</code> and <code>FBIOGATTR ioctl()</code> requests allow access to special features of newer frame buffers. They use the <code>fbsetattr</code> and <code>fbgetattr</code> structures.</p> <p>Some color frame buffers support the <code>FBIOPUTCMAP</code> and <code>FBIOGETCMAP ioctl()</code> requests, which provide access to the colormap. They use the <code>fbcmmap</code> structure.</p> <p>Also, some framebuffer with multiple colormaps will either encode the colormap identifier in the high-order bits of the "index" field in the <code>fbcmmap</code> structure, or use the <code>FBIOPUTCMAPI</code> and <code>FBIOGETCMAPI ioctl()</code> requests.</p> <p><code>FBIOVERTICAL</code> is used to wait for the start of the next vertical retrace period.</p> <p><code>FBIOVRTOFFSET</code> Returns the offset to a read-only <i>vertical retrace page</i> for those framebuffer that support it. This vertical retrace page may be mapped into user space with <code>mmap(2)</code>. The first word of the vertical retrace page (type <code>unsigned int</code>) is a counter that is incremented every time there is a vertical retrace. The user process can use this counter in a variety of ways.</p> <p><code>FBIOMONINFO</code> returns a <code>mon_info</code> structure which contains information about the monitor attached to the framebuffer, if available.</p> <p><code>FBIOSCUSOR</code>, <code>FBIOGCUSOR</code>, <code>FBIOSCURPOS</code> and <code>FBIOGCURPOS</code> are used to control the hardware cursor for those framebuffer that have this feature. <code>FBIOGCURMAX</code> returns the maximum sized cursor supported by the framebuffer. Attempts to create a cursor larger than this will fail.</p> <p>Finally <code>FBIOSDEVINFO</code> and <code>FBIOGDEVINFO</code> are used to transfer variable-length, device-specific information into and out of framebuffer.</p>
SEE ALSO	<code>ioctl(2)</code> , <code>mmap(2)</code> , <code>cgsix(7D)</code>

BUGS The `FBIOSATTR` and `FBIOGATTR` `ioctl()` requests are only supported by frame buffers which emulate older frame buffer types. If a frame buffer emulates another frame buffer, `FBIOGTYPE` returns the emulated type. To get the real type, use `FBIOGATTR`.

The `FBIOGCURPOS` `ioctl` was incorrectly defined in previous operating systems, and older code running in binary compatibility mode may get incorrect results.

fbt(7D)

NAME fbt – DTrace function boundary tracing provider

DESCRIPTION The fbt driver is a DTrace dynamic tracing provider that performs dynamic instrumentation at function boundaries in the Solaris kernel.

The function is the fundamental unit of program text. In a well-designed system, the function performs a discrete and well-defined operation on a specified object or series of like objects. Most functions are implemented by themselves calling functions on encapsulated objects, but some functions —so-called "leaf functions" — are implemented without making further function calls. The Function Boundary Tracing fbt provider contains a mechanism for instrumenting the vast majority of functions in the kernel and offering the instrumentation as a set of DTrace probes.

The fbt driver is not a public interface and you access the instrumentation offered by this provider through DTrace. Refer to the *Solaris Dynamic Tracing Guide* for a description of the public documented interfaces available for the DTrace facility and the probes offered by the fbt provider.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdtrp
Interface Stability	Private

SEE ALSO `dtrace(1M)`, `attributes(5)`, `dtrace(7D)`

Solaris Dynamic Tracing Guide

NAME	fcip – IP/ARP over Fibre Channel datagram encapsulation driver
SYNOPSIS	<code>/dev/fcip</code>
DESCRIPTION	<p>The <code>fcip</code> driver is a Fibre Channel upper layer protocol module for encapsulating IP (IPv4) and ARP datagrams over Fibre Channel. The <code>fcip</code> driver is a loadable, clonable, STREAMS driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code> over any Sun Fibre Channel transport layer-compliant host adapter.</p> <p>The <code>fcip</code> driver complies with the <i>RFC 2625</i> specification for encapsulating IP/ARP datagrams over Fibre Channel, and allows encapsulation of IPv4 only, as specified in <i>RFC 2625</i>. The <code>fcip</code> driver interfaces with the <code>fp(7D)</code> Sun Fibre Channel port driver.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device <code>/dev/fcip</code> is used to access all Fibre Channel ports capable of supporting IP/ARP traffic on the system.</p>
fcip and DLPI	<p>The <code>fcip</code> driver is a "style 2" Data Link Service Provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code> type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <code><sys/dlpi.h></code>. Refer to <code>dlpi(7P)</code> for more information on DLPI primitives.</p> <p>An explicit <code>DL_ATTACH_REQ</code> message must be sent to associate the opened stream with a particular Fibre Channel port (<code>ppa</code>). The <code>ppa</code> ID is interpreted as an unsigned long data type and indicates the corresponding Fibre Channel port driver instance number. An error (<code>DL_ERROR_ACK</code>) is returned by the driver if the <code>ppa</code> field value does not correspond to a valid port driver instance number or if the Fibre Channel port is not <code>ONLINE</code>. Refer to <code>fp(7D)</code> for more details on the Fibre Channel port driver.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to a <code>DL_INFO_REQ</code> from the user are as follows:</p> <ul style="list-style-type: none"> ■ Maximum SDU is 65280 (defined in <i>RFC 2625</i>). ■ Minimum SDU is 0. ■ DLSAP address length is 8. ■ MAC type is <code>DL_ETHER</code>. ■ SAP length is -2. ■ Service mode is <code>DL_CLDLS</code>. ■ Optional quality of service (QOS) fields are set to 0. ■ Provider style is <code>DL_STYLE2</code>. ■ Provider version is <code>DL_VERSION_2</code>. ■ Broadcast address value is <code>0xFFFFFFFF</code>. <p>Once in <code>DL_ATTACHED</code> state, the user must send a <code>DL_BIND_REQ</code> to associate a particular SAP (Service Access Point) with the stream. The <code>fcip</code> driver DLSAP address format consists of the 6-byte physical address component followed immediately by the 2-byte SAP component producing an 8-byte DLSAP address. Applications should not be programmed to use this implementation-specific DLSAP address format, but use information returned in the <code>DL_INFO_ACK</code> primitive to compose and decompose DLSAP addresses. The SAP length, full DLSAP length, and</p>

fcip(7D)

SAP/physical ordering are included within the DL_INFO_ACK. The physical address length is the full DLSAP address length minus the SAP length. The physical address length can also be computed by issuing the DL_PHYS_ADDR_REQ primitive to obtain the current physical address associated with the stream.

Once in the DL_BOUND state, the user can transmit frames on the fibre by sending DL_UNITDATA_REQ messages to the fcip driver. The fcip driver will route received frames up any of the open and bound streams having a SAP which matches the received frame's SAP type as DL_UNITDATA_IND messages. Received Fibre Channel frames are duplicated and routed up multiple open streams if necessary. The DLSAP address contained within the DL_UNITDATA_REQ and DL_UNITDATA_IND messages consists of both the SAP (type) and physical address (WorldWideName) components.

Other Primitives

In Fibre Channel, *multicasting* is defined as an optional service for Fibre Channel classes three and six only. If required, the Fibre Channel broadcast service can be used for multicasting. The RFC 2625 specification does not support IP multicasting or promiscuous mode.

fcip Fibre Channel ELS

The fcip driver will use the FARP Fibre Channel Extended Link Service (ELS), where supported, to resolve WorldWide Names (MAC address) to FC Port Identifiers(Port_ID). The fcip driver also supports InARP to resolve WorldWide Name and Port_ID to an IP address.

FILES

/dev/fcip
fcip character-special device

/kernel/drv/fcip
32-bit ELF kernel driver (x86)

/kernel/drv/amd64/fcip
64-bit ELF kernel driver (x86)

/kernel/drv/sparcv9/fcip
64-bit ELF kernel driver (SPARC)

/kernel/drv/fcip.conf
fcip driver configuration file

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Availability	SUNWfcip

SEE ALSO

netstat(1M), prtconf(1M), driver.conf(4), fp(7D), dlpi(7P)

Writing Device Drivers

IP and ARP over Fibre Channel, RFC 2625 M. Rajagopal, R. Bhagwat, W. Rickard. Gadzoox Networks, June 1999

ANSI X3.230-1994, Fibre Channel Physical and Signalling Interface (FC-PH)

ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL)

NOTES If you use a Fibre Channel adapter with two or more ports that each share a common Node WorldWideName, the `fcip` driver will likely attach to the first port on the adapter.

RFC 2625 requires that both source and destination WorldWideNames have their 4 bit NAA identifiers set to binary '0001,' indicating that an IEEE 48-bit MAC address is contained in the lower 48 bits of the network address fields. For additional details, see the *RFC 2625* specification.

fcplib(7D)

NAME fcp – Fibre Channel protocol driver

DESCRIPTION The fcp driver is the upper layer protocol that supports mechanisms for transporting SCSI-3 commands over Fibre Channel. The fcp driver, which interfaces with the Sun Fibre Channel transport library fctl(7D), supports the standard functions provided by the SCSA interface.

FILES /kernel/drv/fcp
32-bit ELF kernel driver (x86)

/kernel/drv/amd64/fcp
64-bit ELF kernel driver (x86)

/kernel/drv/sparcv9/fcp
64-bit ELF kernel driver (SPARC)

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Interface stability	Unknown
Availability	SUNWfcp

SEE ALSO prtconf(1M), driver.conf(4), fctl(7D), fp(7D), usoc(7D)

Writing Device Drivers

Fibre Channel Physical and Signaling Interface (FC-PH) ANSI X3.230: 1994

Fibre Channel Generic Services (FC-GS-2) Project 1134-D

Fibre Channel Arbitrated Loop (FC-AL) ANSI X3.272-1996

Fibre Channel Protocol for SCSI (FCP) ANSI X3.269-1996

SCSI-3 Architecture Model (SAM) Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA) ANSI X3.270-1996

Fabric Loop Attachment (FC-FLA), NCITS TR-20:1998

- NAME** fctl – Sun Fibre Channel transport library
- DESCRIPTION** The `fctl` kernel module interfaces the Sun Fibre Channel upper layer protocol (ULP) mapping modules with Sun Fibre Channel adapter (FCA) drivers. There are no user-configurable options for this module.
- FILES**
- `/kernel/misc/fctl`
32-bit ELF kernel module (x86)
 - `/kernel/misc/amd64/fctl`
64-bit ELF kernel module (x86)
 - `/kernel/misc/sparcv9/fctl`
64-bit ELF kernel module (SPARC)
- ATTRIBUTES** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Interface stability	Unknown
Availability	SUNWfctl

SEE ALSO [fp\(7D\)](#)

fd(7D)

NAME	fd, fdc – drivers for floppy disks and floppy disk controllers
SYNOPSIS	
SPARC	/dev/diskette0 /dev/rdiskette0
x86	/dev/diskette[0-1] /dev/rdiskette[0-1]
DESCRIPTION	<p>The fd and fdc drivers provide the interfaces to floppy disks using the Intel 8272, Intel 82077, NEC 765, or compatible disk controllers on x86 based systems.</p> <p>The default partitions for the floppy driver are:</p> <ul style="list-style-type: none">a All cylinders except the lastb Only the last cylinderc Entire diskette <p>The fd driver autosenses the density of the diskette.</p> <p>When the floppy is first opened the driver looks for a SunOS label in logical block 0 of the diskette. If attempts to read the SunOS label fail, the open will fail. If block 0 is read successfully but a SunOS label is not found, auto-sensed geometry and default partitioning are assumed.</p> <p>The fd driver supports both block and “raw” interfaces.</p> <p>The block files (/dev/diskette*) access the diskette using the system’s normal buffering mechanism and may be read and written without regard to physical diskette records.</p> <p>There is also a “raw” (/dev/rdiskette*) interface that provides for direct transmission between the diskette and the user’s read or write buffer. A single read(2) or write(2) call usually results in one I/O operation; therefore raw I/O is considerably more efficient when larger blocking factors are used. A blocking factor of no less than 8 Kbytes is recommended. See the Notes section, below, for information on the number of sectors per track.</p>
3.5" Diskettes	For 3.5" double-sided diskettes, the following densities are supported:
SPARC	1.7 Mbyte density 80 cylinders, 21 sectors per track, 1.7 Mbyte capacity high density 80 cylinders, 18 sectors per track, 1.44 Mbyte capacity double density 80 cylinders, 9 sectors per track, 720 Kbyte capacity
x86	extended density 80 cylinders, 36 sectors per track, 2.88 Mbyte capacity 1.7 Mbyte density 80 cylinders, 21 sectors per track, 1.7 Mbyte capacity

	high density	80 cylinders, 18 sectors per track, 1.44 Mbyte capacity
	double density	80 cylinders, 9 sectors per track, 760 Kbyte capacity
5.25" Diskettes	For 5.25" double-sided diskettes on x86 platforms, the densities listed below are supported:	
SPARC	5.25" diskettes are not supported on SPARC platforms.	
x86	high density	80 cylinders, 15 sectors per track, 1.2 Mbyte capacity
	double density	40 cylinders, 9 sectors per track, 360 Kbyte capacity
	double density	40 cylinders, 8 sectors per track, 320 Kbyte capacity
	quad density	80 cylinders, 9 sectors per track, 720 Kbyte capacity
	double density	40 cylinders, 16 sectors per track (256 bytes per sector), 320 Kbyte capacity
	double density	40 cylinders, 4 sectors per track (1024 bytes per sector), 320 Kbyte capacity
ERRORS	EBUSY	During opening, the partition has been opened for exclusive access and another process wants to open the partition. Once open, this error is returned if the floppy disk driver attempted to pass a command to the floppy disk controller when the controller was busy handling another command. In this case, the application should try the operation again.
	EFAULT	An invalid address was specified in an ioctl command (see <code>fdio(7I)</code>).
	EINVAL	The number of bytes read or written is not a multiple of the diskette's sector size. This error is also returned when an unsupported command is specified using the <code>FDIOCMD</code> ioctl command (see <code>fdio(7I)</code>).
	EIO	During opening, the diskette does not have a label or there is no diskette in the drive. Once open, this error is returned if the requested I/O transfer could not be completed.
	ENOSPC	An attempt was made to write past the end of the diskette.
	ENOTTY	The floppy disk driver does not support the requested ioctl functions (see <code>fdio(7I)</code>).
	ENXIO	The floppy disk device does not exist or the device is not ready.
	EROFS	The floppy disk device is opened for write access and the diskette in the drive is write protected.

fd(7D)

x86 Only	ENOSYS	The floppy disk device does not support the requested ioctl function (FDEJECT).
x86 CONFIGURATION	The driver attempts to initialize itself using the information found in the configuration file, /platform/i86pc/kernel/drv/fd.conf. name="fd" parent="fdc" unit=0; name="fd" parent="fdc" unit=1;	
FILES		
SPARC	/platform/sun4u/kernel/drv/fd	driver module
	/usr/include/sys/fdreg.h	structs and definitions for Intel 82072 and 82077 controllers
	/usr/include/sys/fdvar.h	structs and definitions for floppy drivers
	/dev/diskette	device file
	/dev/diskette0	device file
	/dev/rdiskette	raw device file
	/dev/rdiskette0	raw device file
	For ucb Compatibility	
	/dev/fd0 [a-c]	block file
	/dev/rfd0 [a-c]	raw file
	/vol/dev/diskette0	directory containing volume management character device file
	/vol/dev/rdiskette0	directory containing the volume management raw character device file
	/vol/dev/aliases/floppy0	symbolic link to the entry in /vol/dev/rdiskette0
x86	/platform/i86pc/kernel/drv/fd	driver module
	/platform/i86pc/kernel/drv/fd.conf	configuration file for floppy driver
	/platform/i86pc/kernel/drv/fdc	floppy-controller driver module
	/platform/i86pc/kernel/drv/fdc.conf	configuration file for the floppy-controller
	/usr/include/sys/fdc.h	structs and definitions for x86 floppy devices
	/usr/include/sys/fdmedia.h	structs and definitions for x86 floppy media
	x86 First Drive	

/dev/diskette	device file
/dev/diskette0	device file
/dev/rdiskette	raw device file
/dev/rdiskette0	raw device file
For ucb Compatibility	
/dev/fd0 [a-c]	block file
/dev/rfd0 [a-c]	raw file
/vol/dev/diskette0	directory containing volume management character device file
/vol/dev/rdiskette0	directory containing the volume management raw character device file
/vol/dev/aliases/floppy0	symbolic link to the entry in /vol/dev/rdiskette0
x86 Second Drive	
/dev/diskette1	device file
/dev/rdiskette1	raw device file
For ucb Compatibility	
/dev/fd1 [a-c]	block file
/dev/rfd1 [a-c]	raw file
/vol/dev/diskette1	directory containing volume management character device file
/vol/dev/rdiskette1	directory containing the volume management raw character device file
/vol/dev/aliases/floppy1	symbolic link to the entry in /vol/dev/rdiskette1

SEE ALSO fdformat(1), dd(1M), drvconfig(1M), vold(1M), read(2), write(2), driver.conf(4), dkio(7I) fdio(7I)

DIAGNOSTICS

All Platforms

fd<n>: <command name> failed (<sr1> <sr2> <sr3>)

The <command name> failed after several retries on drive <n>. The three hex values in parenthesis are the contents of status register 0, status register 1, and status register 2 of the Intel 8272, the Intel 82072, and the Intel 82077 Floppy Disk Controller on completion of the command, as documented in the data sheet for that part. This error message is usually followed by one of the following, interpreting the bits of the status register:

fd(7D)

```
fd<n>:          not writable
fd<n>:          crc error blk <block number>
                There was a data error on <block number>.
fd<n>:          bad format
fd<n>:          timeout
fd<n>:          drive not ready
fd<n>:          unformatted diskette or no diskette in drive
fd<n>:          block <block number> is past the end!
                (nblk=<total number of blocks>)
                The operation tried to access a block number that is greater than
                the total number of blocks.
fd<n>:          b_bcount 0x<op_size> not % 0x<sect_size>
                The size of an operation is not a multiple of the sector size.
fd<n>:          overrun/underrun
fd<n>:          host bus error. There was a hardware error on a system
                bus.
```

SPARC Only Overrun/underrun errors occur when accessing a diskette while the system is heavily loaded. Decrease the load on the system and retry the diskette access.

NOTES 3.5" high density diskettes have 18 sectors per track and 5.25" high density diskettes have 15 sectors per track. They can cross a track (though not a cylinder) boundary without losing data, so when using `dd(1M)` or `read(2)/write(2)` calls to or from the "raw" diskette, you should specify `bs=18k` or multiples thereof for 3.5" diskettes, and `bs=15k` or multiples thereof for 5.25" diskettes.

The SPARC `fd` driver is *not* an unloadable module.

Under Solaris (x86 Edition), the configuration of the floppy drives is specified in CMOS configuration memory. Use the BIOS setup program for the system to define the diskette size and density/capacity for each installed drive. Note that MS-DOS may operate the floppy drives correctly, even though the CMOS configuration may be in error. Solaris (x86 Edition) relies on the CMOS configuration to be accurate.

NAME fdio – floppy disk control operations

SYNOPSIS `#include <sys/fdio.h>`

DESCRIPTION The Solaris floppy driver supports a set of `ioctl(2)` requests for getting and setting the floppy drive characteristics. Basic to these `ioctl()` requests are the definitions in `<sys/fdio.h>`.

IOCTLS The following `ioctl()` requests are available on the Solaris floppy driver.

FDDEFGECHAR x86 based systems: This `ioctl()` forces the floppy driver to restore the diskette and drive characteristics and geometry, and partition information to default values based on the device configuration.

FDGETCHANGE The argument is a pointer to an `int`. This `ioctl()` returns the status of the diskette-changed signal from the floppy interface. The following defines are provided for cohesion.

Note: For x86 based systems, use `FDGC_DETECTED` (which is available only on x86 based systems) instead of `FDGC_HISTORY`.

```

/*
 * Used by FDGETCHANGE, returned state of the sense disk change bit.
 */
#define FDGC_HISTORY 0x01 /* disk has changed since insertion or
                          last FDGETCHANGE call */
#define FDGC_CURRENT 0x02 /* if set, indicates drive has floppy,
                          > otherwise, drive is empty */
#define FDGC_CURWPROT 0x10 /* current state of write protect */
#define FDGC_DETECTED 0x20 /* previous state of DISK CHANGE */

```

FDIOGCHAR The argument is a pointer to an `fd_char` structure (described below). This `ioctl()` gets the characteristics of the floppy diskette from the floppy controller.

FDIOSCHAR The argument is a pointer to an `fd_char` structure (described below). This `ioctl()` sets the characteristics of the floppy diskette for the floppy controller. Typical values in the `fd_char` structure for a high density diskette:

```

field value
fdc_medium      0
fdc_transfer_rate 500
fdc_ncyl        80
fdc_nhead       2
fdc_sec_size     512
fdc_secptrack   18
fdc_steps       -1 { This field doesn't apply. }

```

```

/*
 * Floppy characteristics
 */
struct fd_char {

```

fdio(7I)

```

uchar_t fdc_medium;      /* equals 1 if floppy is medium density format */
int fdc_transfer_rate;   /* transfer rate */
int fdc_ncyl;            /* number of cylinders */
int fdc_nhead;          /* number of heads */
int fdc_sec_size;       /* sector size */
int fdc_secptrack;      /* sectors per track */
int fdc_steps;          /* no. of steps per data track */
};

```

FDGETDRIVECHAR The argument to this `ioctl()` is a pointer to an `fd_drive` structure (described below). This `ioctl()` gets the characteristics of the floppy drive from the floppy controller.

FDSETDRIVECHAR x86 based systems: The argument to this `ioctl()` is a pointer to an `fd_drive` structure (described below). This `ioctl()` sets the characteristics of the floppy drive for the floppy controller. Only `fdd_steprate`, `fdd_headsettle`, `fdd_motoron`, and `fdd_motoroff` are actually used by the floppy disk driver.

```

/*
 * Floppy Drive characteristics
 */
struct fd_drive {
    int fdd_ejectable;      /* does the drive support eject? */
    int fdd_maxsearch;     /* size of per-unit search table */
    int fdd_writeprecomp;  /* cyl to start write precompensation */
    int fdd_writereduce;   /* cyl to start recucing write current */
    int fdd_stepwidth;     /* width of step pulse in 1 us units */
    int fdd_steprate;      /* step rate in 100 us units */
    int fdd_headsettle;    /* delay, in 100 us units */
    int fdd_headload;     /* delay, in 100 us units */
    int fdd_headunload;    /* delay, in 100 us units */
    int fdd_motoron;       /* delay, in 100 ms units */
    int fdd_motoroff;      /* delay, in 100 ms units */
    int fdd_precomplevel;  /* bit shift, in nano-secs */
    int fdd_pins;          /* defines meaning of pin 1, 2, 4 and 34 */
    int fdd_flags;         /* TRUE READY, Starting Sector #, & Motor On */
};

```

FDGETSEARCH Not available.

FDSETSEARCH Not available.

FDEJECT SPARC: This `ioctl()` requests the floppy drive to eject the diskette.

FDIOCMD The argument is a pointer to an `fd_cmd` structure (described below). This `ioctl()` allows access to the floppy diskette using the floppy device driver. Only the `FDCMD_WRITE`, `FDCMD_READ`, and `FDCMD_FORMAT_TRACK` commands are currently available.

```

struct fd_cmd {
    ushort_t fdc_cmd;      /* command to be executed */
    int      fdc_flags;   /* execution flags (x86 only) */
    daddr_t  fdc_blkno;   /* disk address for command */
    int      fdc_secnt;   /* sector count for command */
    caddr_t  fdc_bufaddr; /* user's buffer address */
    uint_t   fdc_buflen;  /* size of user's buffer */
};

```

Please note that the `fdc_buflen` field is currently unused. The `fdc_secnt` field is used to calculate the transfer size, and the buffer is assumed to be large enough to accommodate the transfer.

```

{
/*
 * Floppy commands
 */
#define FDCMD_WRITE    1
#define FDCMD_READ     2
#define FDCMD_SEEK    3
#define FDCMD_REZERO   4
#define FDCMD_FORMAT_UNIT  5
#define FDCMD_FORMAT_TRACK 6
};

```

FDRAW The argument is a pointer to an `fd_raw` structure (described below). This `ioctl()` allows direct control of the floppy drive using the floppy controller. Refer to the appropriate floppy-controller data sheet for full details on required command bytes and returned result bytes. The following commands are supported.

```

/*
 * Floppy raw commands
 */
#define FDRAW_SPECIFY    0x03
#define FDRAW_READID    0x0a    (x86 only)
#define FDRAW_SENSE_DRV  0x04
#define FDRAW_REZERO    0x07
#define FDRAW_SEEK      0x0f
#define FDRAW_SENSE_INT  0x08    (x86 only)
#define FDRAW_FORMAT    0x0d
#define FDRAW_READTRACK 0x02
#define FDRAW_WRCMD     0x05
#define FDRAW_RDCMD     0x06
#define FDRAW_WRITEDEL  0x09
#define FDRAW_READDEL   0x0c

```

Please note that when using `FDRAW_SEEK` or `FDRAW_REZERO`, the driver automatically issues a `FDRAW_SENSE_INT` command to clear the interrupt from the `FDRAW_SEEK` or the `FDRAW_REZERO`. The result bytes returned by these commands are the results from the `FDRAW_SENSE_INT` command. Please see the floppy-controller data sheet for more details on `FDRAW_SENSE_INT`.

fdio(7I)

```
/*
 * Used by FDRAW
 */
struct    fd_raw {
char      fdr_cmd[10];      /* user-supplied command bytes */
short     fdr_cnum;        /* number of command bytes */
char      fdr_result[10];  /* controller-supplied result bytes */
ushort_t  fdr_nbytes;      /* number to transfer if read/write command */
char      *fdr_addr;       /* where to transfer if read/write command */
};
```

SEE ALSO [ioctl\(2\)](#), [dkio\(7I\)](#), [fd\(7D\)](#), [hdio\(7I\)](#)

NAME	ffb – 24-bit UPA color frame buffer and graphics accelerator					
DESCRIPTION	<p>ffb is a 24-bit UPA-based color frame buffer and graphics accelerator which comes in the two configurations: single buffered frame and double buffered frame.</p> <table> <tr> <td>Single buffered frame buffer</td> <td>Consists of 32 video memory planes of 1280 x 1024 pixels, including 24-bit single-buffering and 8-bit X planes.</td> </tr> <tr> <td>Double buffered frame buffer</td> <td>Consists of 96 video memory planes of 1280 x 1024 pixels, including 24-bit double-buffering, 8-bit X planes, 28-bit Z-buffer planes and 4-bit Y planes.</td> </tr> </table> <p>The driver supports the following frame buffer ioctls which are defined in fbio(7I):</p> <pre>FBIOPUTCMAP, FBIOGETCMAP, FBIOSVIDEO, FBIOGVIDEO, FBIOVERTICAL, FBIOSCURSOR, FBIOGCURSOR, FBIOSCURPOS, FBIOGCURPOS, FBIOGCURMAX, FBIO_WID_PUT, FBIO_WID_GET</pre>		Single buffered frame buffer	Consists of 32 video memory planes of 1280 x 1024 pixels, including 24-bit single-buffering and 8-bit X planes.	Double buffered frame buffer	Consists of 96 video memory planes of 1280 x 1024 pixels, including 24-bit double-buffering, 8-bit X planes, 28-bit Z-buffer planes and 4-bit Y planes.
Single buffered frame buffer	Consists of 32 video memory planes of 1280 x 1024 pixels, including 24-bit single-buffering and 8-bit X planes.					
Double buffered frame buffer	Consists of 96 video memory planes of 1280 x 1024 pixels, including 24-bit double-buffering, 8-bit X planes, 28-bit Z-buffer planes and 4-bit Y planes.					
FILES	/dev/fbs/ffb0 device special file					
SEE ALSO	ffbconfig(1M), mmap(2), fbio(7I)					

flowacct(7ipp)

NAME flowacct – Flow Accounting module

DESCRIPTION The flow accounting module `flowacct` enables you to record flow details. You use flow details to gather statistics and/or for billing purposes. Accounting consists of recording flow details in a location you designate and in a format that you can retrieve at a later stage. IPQoS accounting relies on the exact mechanism to store and retrieve flow information.

A flow is defined by the 5-tuple - `saddr`, `sport`, `daddr`, `dport` and protocol.

Typically, the accounting module is the last datapath element in a sequence of actions. Flow attributes include ToS/DS, user id, project id, creation time (time the flow was created), last seen (when pkts for the flow were last seen), action name (instance that recorded the flow information), nbytes and npackets. Attributes are split into groups entitled *basic* and *extended*. The basic group records only the nbytes, npackets and action name, while the extended group is a superset of the basic group and records all attributes. The attributes to be recorded, in addition to the accounting file that contains flow details, are selected using `acctadm(1M)`. The `flowacct` module does not provide a mechanism to retrieve flow information from the accounting file nor to interpret the retrieved information.

STATISTICS The flowacct module exports the following statistics available through `kstat`:

```

module: flowacct           instance: <action id>
  name: Flowacct statistics class <action name>
    bytes_in_tbl           <bytes in the flow table>
    epackets               <packets in error>
    flows_in_tbl           <flow records in the flow table>
    nbytes                 <number of bytes through this instance>
    npackets               <number of packets>
    usedmem                <memory, in bytes, used by the flow table>
  
```

FILES /kernel/ipp/sparcv9/flowacct
64-bit module (SPARC only.)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWqos

SEE ALSO `ipqosconf(1M)`, `acctadm(1M)`, `libexacct3LIB`, `dlcosmk(7IPP)`, `dscpmk(7IPP)`, `ipqos(7IPP)`, `igppc(7IPP)`, `tokenmt(7IPP)`, `tswtclmt(7IPP)`

NAME	fp – Sun Fibre Channel port driver
DESCRIPTION	<p>The fp driver is a Sun fibre channel nexus driver that enables fibre channel topology discovery, device discovery, fibre channel adapter port management and other capabilities through well-defined fibre channel adapter driver interfaces.</p> <p>The fp driver requires the presence of a fabric name server in fabric and public loop topologies to discover fibre channel devices. In private loop topologies, the driver discovers devices by performing PLOGI to all valid AL_PAs, provided that devices do not participate in LIRP and LILP stages of loop initialization.</p>
CONFIGURATION	<p>The fp driver is configured by defining properties in the fp.conf file. The fp driver supports the following properties:</p>
mpxio-disable	<p>Solaris I/O multipathing is enabled or disabled on fibre channel devices with the mpxio-disable property. Specifying mpxio-disable="no" activates I/O multipathing, while mpxio-disable="yes" disables the feature.</p> <p>Solaris I/O multipathing may be enabled or disabled on a per port basis. Per port settings override the global setting for the specified ports. The following example shows how to disable multipathing on port 0 whose parent is /pci@8,600000/SUNW,qlc@4:</p> <pre>name="fp" parent="/pci@8,600000/SUNW,qlc@4" port=0 mpxio-disable="yes";</pre>
manual_configuration_only	<p>Automatic configuration of SCSI devices in the fabric is enabled by default and thus allows all devices discovered in the SAN zone to be enumerated in the kernel's device tree automatically. The manual_configuration_only property may be configured to disable the default behavior and force the manual configuration of the devices in the SAN. Specifying manual_configuration_only=1 disables the automatic configuration of devices.</p>
FILES	<p>/kernel/drv/fp 32-bit ELF kernel driver (x86)</p> <p>/kernel/drv/amd64/fp 64-bit ELF kernel driver (x86)</p> <p>/kernel/drv/sparcv9/fp 64-bit ELF kernel driver (SPARC)</p> <p>/kernel/drv/fp.conf fp driver configuration file.</p>
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p>

ATTRIBUTE TYPE	ATTRIBUTE VALUE
mpxio-disable	Unstable

fp(7d)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
manual_configuration_only	Obsolete
Availability	SUNWfctl

SEE ALSO `cfgadm_fp(1M)`, `prtconf(1M)`, `stmsboot(1M)`, `driver.conf(4)`, `attributes(5)`, `fcp(7D)`, `fctl(7D)`, `scsi_vhci(7D)`

Writing Device Drivers

Fibre Channel Physical and Signaling Interface (FC-PH) ANSI X3.230: 1994

Fibre Channel Generic Services (FC-GS-2) Project 1134-D

Fibre Channel Arbitrated Loop (FC-AL) ANSI X3.272-1996

Fibre Channel Protocol for SCSI (FCP) ANSI X3.269-1996

SCSI-3 Architecture Model (SAM) Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA) ANSI X3.270-1996

SCSI Direct Attach (FC-PLDA) ANSI X3.270-1996

SCSI Direct Attach (FC-PLDA) NCITS TR-19:1998

Fabric Loop Attachment (FC-FLA), NCITS TR-20:1998

NAME	FSS – Fair share scheduler
DESCRIPTION	<p>The fair share scheduler (FSS) guarantees application performance by explicitly allocating shares of CPU resources to projects. A share indicates a project's entitlement to available CPU resources. Because shares are meaningful only in comparison with other project's shares, the absolute quantity of shares is not important. Any number that is in proportion with the desired CPU entitlement can be used.</p> <p>The goals of the FSS scheduler differ from the traditional time-sharing scheduling class (TS). In addition to scheduling individual LWPs, the FSS scheduler schedules projects against each other, making it impossible for any project to acquire more CPU cycles simply by running more processes concurrently.</p> <p>A project's entitlement is individually calculated by FSS independently for each processor set if the project contains processes bound to them. If a project is running on more than one processor set, it can have different entitlements on every set. A project's entitlement is defined as a ratio between the number of shares given to a project and the sum of shares of all active projects running on the same processor set. An active project is one that has at least one running or runnable process. Entitlements are recomputed whenever any project becomes active or inactive, or whenever the number of shares is changed.</p> <p>Processor sets represent virtual machines in the FSS scheduling class and processes are scheduled independently in each processor set. That is, processes compete with each other only if they are running on the same processor set. When a processor set is destroyed, all processes that were bound to it are moved to the default processor set, which always exists. Empty processor sets (that is, sets without processors in them) have no impact on the FSS scheduler behavior.</p> <p>If a processor set contains a mix of TS/IA and FSS processes, the fairness of the FSS scheduling class can be compromised because these classes use the same range of priorities. Fairness is most significantly affected if processes running in the TS scheduling class are CPU-intensive and are bound to processors within the processor set. As a result, you should avoid having processes from TS/IA and FSS classes share the same processor set. RT and FSS processes use disjoint priority ranges and therefore can share processor sets.</p> <p>As projects execute, their CPU usage is accumulated over time. The FSS scheduler periodically decays CPU usages of every project by multiplying it with a decay factor, ensuring that more recent CPU usage has greater weight when taken into account for scheduling. The FSS scheduler continually adjusts priorities of all processes to make each project's relative CPU usage converge with its entitlement.</p> <p>While FSS is designed to fairly allocate cycles over a long-term time period, it is possible that projects will not receive their allocated shares worth of CPU cycles due to uneven demand. This makes one-shot, instantaneous analysis of FSS performance data unreliable.</p> <p>Note that share is not the same as utilization. A project may be allocated 50% of the system, although on the average, it uses just 20%. Shares serve to cap a project's CPU usage only when there is competition from other projects running on the same</p>

processor set. When there is no competition, utilization may be larger than entitlement based on shares. Allocating a small share to a busy project slows it down but does not prevent it from completing its work if the system is not saturated.

The configuration of CPU shares is managed by the name server as a property of the `project(4)` database. In the following example, an entry in the `/etc/project` file sets the number of shares for project "x-files" to 10:

```
x-files:100:::project.cpu-shares=(privileged,10,none)
```

Projects with undefined number of shares are given one share each. This means that such projects are treated with equal importance. Projects with 0 shares only run when there are no projects with non-zero shares competing for the same processor set. The maximum number of shares that can be assigned to one project is 65535.

You can use the `prctl(1)` command to determine the current share assignment for a given project:

```
$ prctl -n project.cpu-shares -i project x-files
```

or to change the amount of shares if you have root privileges:

```
# prctl -r -n project.cpu-shares -v 5 -i project x-files
```

See the `prctl(1)` man page for additional information on how to modify and examine resource controls associated with active processes, tasks, or projects on the system. See `resource_controls(5)` for a description of the resource controls supported in the current release of the Solaris operating system.

By default, project "system" (project ID 0) includes all system daemons started by initialization scripts and has an "unlimited" amount of shares. That is, it is always scheduled first no matter how many shares are given to other projects.

The following command sets FSS as the default scheduler for the system:

```
# dispadmin -d FSS
```

This change will take effect on the next reboot. Alternatively, you can move processes from the time-share scheduling class (as well as the special case of `init`) into the FSS class without changing your default scheduling class and rebooting by becoming root, and then using the `prionctl(1)` command, as shown in the following example:

```
# prionctl -s -c FSS -i class TS
# prionctl -s -c FSS -i pid 1
```

CONFIGURING SCHEDULER WITH DISPADMIN

You can use the `dispadmin(1M)` command to examine and "tune" the FSS scheduler's time quantum value. Time quantum is the amount of time that a thread is allowed to run before it must relinquish the processor. The following example dumps the current time quantum for the fair share scheduler:

```
$ dispadmin -g -c FSS
#
# Fair Share Scheduler Configuration
#
```

```
RES=1000
#
# Time Quantum
#
QUANTUM=110
```

The value of the QUANTUM represents some fraction of a second with the fractional value determined by the reciprocal value of RES. With the default value of RES = 1000, the reciprocal of 1000 is .001, or milliseconds. Thus, by default, the QUANTUM value represents the time quantum in milliseconds.

If you change the RES value using `dispadmin` with the `-r` option, you also change the QUANTUM value. For example, instead of quantum of 110 with RES of 1000, a quantum of 11 with a RES of 100 results. The fractional unit is different while the amount of time is the same.

You can use the `-s` option to change the time quantum value. Note that such changes are not preserved across reboot. Please refer to the `dispadmin(1M)` man page for additional information.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SUNWcsu

SEE ALSO `prctl(1)`, `priocntl(1)`, `dispadmin(1M)`, `psrset(1M)`, `priocntl(2)`, `project(4)`, `attributes(5)`, `resource_controls(5)`

System Administration Guide: N1 Grid Containers, Resource Management, and Solaris Zones

ge(7D)

NAME	ge – GEM Gigabit-Ethernet device driver
SYNOPSIS	/dev/ge
DESCRIPTION	<p>The ge Gigabit-Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code> over GEM, SBus and PCI Gigabit-Ethernet add-in adapters. Multiple GEM-based adapters installed within the system are supported by the driver. The ge driver provides basic support for the GEM-based Ethernet hardware and handles the <code>SUNW, sbus-gem</code> (SBus GEM) and <code>pci108e, 2bad</code> (PCI GEM) devices. Functions include chip initialization, frame transmit and receive, multicast and promiscuous support, and error recovery and reporting.</p> <p>The GEM device provides 1000BASE-SX networking interfaces using the GEM ASIC external SERDES and fiber optical transceiver. The GEM ASIC provides the appropriate bus interface, MAC functions and physical code sub-layer (PCS) functions. The external SERDES connects to a fiber transceiver and provides the physical connection.</p> <p>The 1000Base-SX standard specifies an auto-negotiation protocol to automatically select the mode of operation. In addition to duplex operation, the GEM ASIC can auto-negotiate for <i>IEEE 802.3x</i> frame based flow control capabilities. The GEM PCS is capable of performing auto-negotiation using the remote (or link partner) link end and receives the capabilities of the remote end. It selects the highest common denominator mode of operation based on priorities. The ge driver also supports forced-mode operation under which the driver selects the mode of operation.</p>
APPLICATION PROGRAMMING INTERFACE ge and DLPI	<p>The cloning character-special device <code>/dev/ge</code> is used to access all ge controllers installed within the system.</p> <p>The ge driver is a Style 2 data link service provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code> type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <code><sys/dlpi.h></code>. Refer to <code>dlpi(7P)</code> for more information.</p> <p>You must send an explicit <code>DL_ATTACH_REQ</code> message to associate the opened stream with a particular device (<code>ppa</code>). The <code>ppa</code> ID is interpreted as an unsigned long data type and indicates the corresponding device instance (unit) number. An error (<code>DL_ERROR_ACK</code>) is returned by the driver if the <code>ppa</code> field value does not correspond to a valid device instance number for the system. The device is initialized on first attach and de-initialized (stopped) upon last detach.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> are:</p> <ul style="list-style-type: none">■ Maximum SDU is 1500 (<code>ETHERMTU</code> - defined in <code><sys/ethernet.h></code>).■ Minimum SDU is 0.■ <code>dlsap</code> address length is 8.■ MAC type is <code>DL_ETHER</code>.■ <code>sap</code> length value is -2, meaning the physical address component is followed immediately by a 2 byte <code>sap</code> component within the DLSAP address.

- Service mode is DL_CLDLS.
- Quality of service (QOS) is not supported; accordingly, QOS fields are 0.
- Provider style is DL_STYLE2.
- Version is DL_VERSION_2.
- Broadcast address value is Ethernet/IEEE broadcast address (0xFFFFFFFF).

Once in the DL_ATTACHED state, you must send a DL_BIND_REQ to associate a particular Service Access Pointer (SAP) with the stream. The ge driver interprets the sap field within the DL_BIND_REQ as an Ethernet *type*; accordingly, valid values for the sap field are in the [0-0xFFFF] range. Only one Ethernet type can be bound to the stream at any time.

If you select a sap with a value of 0, the receiver will be in 802.3 mode. All frames received from the media with a type field in the range [0-1500] are assumed to be 802.3 frames and are routed up all open streams bound to sap value 0. If more than one stream is in 802.3 mode, the frame will be duplicated and routed up multiple streams as DL_UNITDATA_IND messages.

In transmission, the driver checks the sap field of the DL_BIND_REQ to determine if the sap value is 0 and the destination type field is in the range [0-1500]. If either is true, the driver computes the length of the message, not including initial M_PROTO mblk (message block), of all subsequent DL_UNITDATA_REQ messages and transmits 802.3 frames of that value in the MAC frame header length field.

The ge driver DLSAP address format consists of the 6 byte physical (Ethernet) address component followed immediately by the 2 byte sap (type) component producing an 8 byte DLSAP address. Applications should *not* hard code to this particular implementation-specific DLSAP address format, but use information returned in the DL_INFO_ACK primitive to compose and decompose DLSAP addresses. The sap length, full DLSAP length and sap physical ordering are included within the DL_INFO_ACK. The physical address length can be computed by subtracting the sap length from the full DLSAP address length or by issuing the DL_PHYS_ADDR_REQ to obtain the current physical address associated with the stream.

Once in the DL_BOUND state, you may transmit frames on the Ethernet by sending DL_UNITDATA_REQ messages to the ge driver. The ge driver will route received Ethernet frames up all open and bound streams having a sap which matches the Ethernet type as DL_UNITDATA_IND messages. Received Ethernet frames are duplicated and routed up multiple open streams if necessary. The DLSAP address contained within the DL_UNITDATA_REQ and DL_UNITDATA_IND messages consists of both the sap (type) and physical (Ethernet) components.

In addition to the mandatory connectionless DLPI message set, the driver additionally supports ge primitives.

ge(7D)

ge Primitives

The `DL_ENABMULTI_REQ` and `DL_DISABMULTI_REQ` primitives enable/disable reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. The `DL_ENABMULTI_REQ` and `DL_DISABMULTI_REQ` primitives are accepted by the driver in any state following `DL_ATTACHED`.

With the `DL_PROMISC_PHYS` flag set in the `dl_level` field, the `DL_PROMISCON_REQ` and `DL_PROMISCOFF_REQ` primitives enable/disable reception of all promiscuous mode frames on the media including frames generated by the local host. When used with the `DL_PROMISC_SAP` flag set, this enables/disables reception of all `sap` (Ethernet type) values. When used with the `DL_PROMISC_MULTI` flag set, this enables/disables reception of all multicast group addresses. The effect of each is always on a per-stream basis and independent of the other `sap` and physical level configurations on the stream or other streams.

The `DL_PHYS_ADDR_REQ` primitive returns the six octet Ethernet address currently associated (attached) to the stream in the `DL_PHYS_ADDR_ACK` primitive. This primitive is valid only in states following a successful `DL_ATTACH_REQ`.

The `DL_SET_PHYS_ADDR_REQ` primitive changes the 6 octet Ethernet address currently associated (attached) to the stream. The credentials of the process which originally opened the stream must be superuser or `EPERM` is returned in the `DL_ERROR_ACK`. The `DL_SET_PHYS_ADDR_REQ` primitive is destructive and affects all other current and future streams attached to this device. A `M_ERROR` is sent up all other streams attached to the device when `DL_SET_PHYS_ADDR_REQ` is successful on the stream. Once changed, all streams subsequently opened and attached to the device will obtain the new physical address. Once changed, the physical address will remain until `DL_SET_PHYS_ADDR_REQ` is used to change the physical address again or the system is rebooted, whichever comes first.

ge DRIVER

By default, the `ge` driver performs auto-negotiation to select the mode and flow control capabilities of the link. The link can be in one of the following modes:

- 1000 Mbps, full-duplex
- 1000 Mbps, half-duplex
- Symmetric pause
- Asymmetric pause

Speeds and modes are described in the 1000Base-TX standard.

The auto-negotiation protocol automatically selects:

- Operation mode (half-duplex or full-duplex)
- Flow control capability (Symmetric and/or Asymmetric)

The auto-negotiation protocol:

- Gets all the modes of operation supported by the link partner.
- Advertises its capabilities to the link partner.
- Selects the highest common denominator mode of operation based on the priorities.

When by default, auto-negotiation is used to bring up the link and select the common mode of operation with the link partner, the GEM hardware is capable of all of the operating modes listed above. The PCS also supports forced-mode of operation under which the driver can select the mode of operation and flow control capabilities using the `ndd` utility.

The GEM device also supports programmable Inter-Packet Gap (IPG) parameters `ipg1` and `ipg2`. By default, the driver sets `ipg1` to 8 byte-times and `ipg2` to 4 byte-times, (the standard values.) You may want to alter these values from the standard 1000 Mbps IPG set to 0.096 microseconds.

ge Parameter List You can use the `ge` driver to set and get parameters for the GEM device. The parameter list includes current transceiver status, current link status, inter-packet gap, PCS capabilities and link partner capabilities.

The PCS has two set of capabilities. One set reflects the capabilities of the hardware and are read-only. The second set are read/write and are used in speed selection and reflect the values you choose. At boot time, both sets will be the same. The link partner capabilities are read only and cannot be modified.

FILES `/dev/ge` `ge` special character device
`/kernel/drv/ge.conf` System wide default device driver properties

SEE ALSO `ndd(1M)`, `netstat(1M)`, `driver.conf(4)`, `dlpi(7P)`, `hme(7D)`, `qfe(7D)`

gld(7D)

NAME	gld – Generic LAN Driver
SYNOPSIS	<pre>#include <sys/stropts.h> #include <sys/stream.h> #include <sys/dlpi.h> #include <sys/gld.h></pre>
INTERFACE LEVEL	Solaris architecture specific (Solaris DDI).
DESCRIPTION	<p>GLD is a multi-threaded, clonable, loadable kernel module providing support for Solaris local area network (LAN) device drivers. LAN drivers in Solaris are STREAMS-based drivers that use the Data Link Provider Interface (DLPI) to communicate with network protocol stacks. These protocol stacks use the network drivers to send and receive packets on a local area network. A network device driver must implement and adhere to the requirements imposed by the DDI/DKI specification, STREAMS specification, DLPI specification, and programmatic interface of the device itself.</p> <p>GLD implements most STREAMS and DLPI functionality required of a Solaris LAN driver. Several Solaris network drivers are implemented using GLD.</p> <p>A Solaris network driver implemented using GLD comprises two distinct parts: a generic component that deals with STREAMS and DLPI interfaces, and a device-specific component that deals with the particular hardware device. The device-specific module indicates its dependency on the GLD module and registers itself with GLD from within the driver's <code>attach(9E)</code> function. Once it is successfully loaded, the driver is DLPI-compliant. The device-specific part of the driver calls <code>gld(9F)</code> functions when it receives data or needs some service from GLD. GLD makes calls into the <code>gld(9E)</code> entry points of the device-specific driver through pointers provided to GLD by the device-specific driver when it registered itself with GLD. The <code>gld_mac_info(9S)</code> structure is the main data interface between GLD and the device-specific driver.</p> <p>The GLD facility currently supports devices of type <code>DL_ETHER</code>, <code>DL_TPR</code>, and <code>DL_FDDI</code>. GLD drivers are expected to process fully-formed MAC-layer packets and should not perform logical link control (LLC) handling.</p> <p>Note – Support for the <code>DL_TPR</code> and <code>DL_FDDI</code> media types in GLD is obsolete and may be removed in a future release of Solaris.</p> <p>In some cases, it may be necessary or desirable to implement a full DLPI-compliant driver without using the GLD facility. This is true for devices that are not IEEE 802-style LAN devices, or where a device type or DLPI service not supported by GLD is required.</p>
Device Naming Constraints	The name of the device-specific driver module must adhere to the naming constraints outlined in the NOTES section of <code>dlpi(7P)</code> .

**Type DL_ETHER:
Ethernet V2 and
ISO 8802-3 (IEEE
802.3)**

For devices designated type DL_ETHER, GLD provides support for both Ethernet V2 and ISO 8802-3 (IEEE 802.3) packet processing. Ethernet V2 enables a data link service user to access and use any of a variety of conforming data link service providers without special knowledge of the provider's protocol. A service access point (SAP) is the point through which the user communicates with the service provider.

SAP 0 denotes that the user wishes to use 802.3 mode. In transmission, GLD checks the destination SAP value of the DL_UNITDATA_REQ and the SAP value to which the stream is bound. If both are 0, the GLD computes the length of the packet payload and transmits 802.3 frames having that length in the MAC frame header type field. Such lengths will never exceed 1500.

All frames received from the media that have a type field in the range [0-1500] are assumed to be 802.3 frames and are routed up all open streams that are in 802.3 mode, (those streams bound to a SAP value in of 0. If more than one stream is in 802.3 mode, the incoming frame is duplicated and routed up each such stream.

Streams bound to a SAP value of 1536 or greater receive incoming packets whose Ethernet MAC header type value exactly matches the value of the SAP to which the stream is bound. SAP values in the range [1-1535] are undefined and should not be used.

**Types DL_TPR
and DL_FDDI:
SNAP Processing**

Note – Support for the DL_TPR and DL_FDDI media types in GLD is obsolete and may be removed in a future release of Solaris.

For media types DL_TPR and DL_FDDI, GLD implements minimal SNAP (Sub-Net Access Protocol) processing for SAP values of 1536 or greater. A SAP value of 0 denotes that the user wishes to use LLC mode. SAP values in the range [1-1535] have undefined semantics and should not be used.

SNAP headers are carried under LLC headers with destination SAP 0xAA. For outgoing packets with SAP values greater than 1535, GLD creates an LLC+SNAP header that always looks like:

“AA AA 03 00 00 00 XX XX”

where “XX XX” represents the 16-bit SAP, corresponding to the Ethernet V2 style “type.” This is the only class of SNAP header that is processed - non-zero OUI fields, and LLC control fields other than 03 are considered to be LLC packets with SAP 0xAA.

A DL_UNITDATA_REQ message specifying a destination SAP value of 0, passed down a stream bound to SAP 0, is assumed to contain an LLC packet and will not undergo SNAP processing.

gld(7D)

**Type DL_TPR:
Source Routing**

Incoming packets are examined to ascertain whether they fall into the format specified above. Packets that do will be passed to streams bound to the packet's 16-bit SNAP type, as well as being passed to any stream in LLC mode (those bound to a SAP value of 0).

Note – Support for the DL_TPR media type in GLD is obsolete and may be removed in a future release of Solaris.

For type DL_TPR devices, GLD implements minimal support for source routing. Source routing enables a station that is sending a packet across a bridged medium to specify (in the packet MAC header) routing information that determines the route that the packet will take through the network.

Functionally, the source routing support provided by GLD learns routes, solicits and responds to requests for information about possible multiple routes and selects among the multiple routes that are available. It adds *Routing Information Fields* to the MAC headers of outgoing packets and recognizes such fields in incoming packets.

GLD's source routing support does not implement the full *Route Determination Entity* (RDE) specified in *ISO 8802-2 (IEEE 802.2)* Section 9. However, it is designed to interoperate with any such implementations that may exist in the same (or a bridged) network.

**Style 1 and 2
Providers**

GLD implements both Style 1 and Style 2 providers. A physical point of attachment (PPA) is the point at which a system attaches itself to a physical communication medium. All communication on that physical medium funnels through the PPA. The Style 1 provider attaches the stream to a particular PPA based on the major/minor device that has been opened. The Style 2 provider requires the DLS user to explicitly identify the desired PPA using DL_ATTACH_REQ. In this case, open(9E) creates a stream between the user and GLD and DL_ATTACH_REQ subsequently associates a particular PPA with that stream. Style 2 is denoted by a minor number of zero. If a device node whose minor number is not zero is opened, Style 1 is indicated and the associated PPA is the minor number minus 1. In both Style 1 and Style 2 opens, the device is cloned.

**Implemented
DLPI Primitives**

GLD implements the following DLPI primitives:

The DL_INFO_REQ primitive requests information about the DLPI stream. The message consists of one M_PROTO message block. GLD returns device-dependent values in the DL_INFO_ACK response to this request, based on information the GLD-based driver specified in the gld_mac_info(9S) structure passed to gld_register(). However GLD returns the following values on behalf of all GLD-based drivers:

- The version is DL_VERSION_2.
- The service mode is DL_CLDLS — GLD implements connectionless-mode service.
- The provider style is DL_STYLE1 or DL_STYLE2, depending on how the stream was opened.

The `DL_ATTACH_REQ` primitive is called to associate a PPA with a stream. This request is needed for Style 2 DLS providers to identify the physical medium over which the communication will transpire. Upon completion, the state changes from `DL_UNATTACHED` to `DL_UNBOUND`. The message consists of one `M_PROTO` message block. This request may not be issued when using the driver in Style 1 mode; streams opened using Style 1 are already attached to a PPA by the time the open completes.

The `DL_DETACH_REQ` primitive requests to detach the PPA from the stream. This is only allowed if the stream was opened using Style 2.

The `DL_BIND_REQ` and `DL_UNBIND_REQ` primitives bind and unbind a DLSAP to the stream. The PPA associated with each stream will have been initialized upon completion of the processing of the `DL_BIND_REQ`. Multiple streams may be bound to the same SAP; each such stream receives a copy of any packets received for that SAP.

The `DL_ENABMULTI_REQ` and `DL_DISABMULTI_REQ` primitives enable and disable reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. The stream must be attached to a PPA for these primitives to be accepted.

The `DL_PROMISCON_REQ` and `DL_PROMISCOFF_REQ` primitives enable and disable promiscuous mode on a per-stream basis, either at a physical level or at the SAP level. The DL Provider will route all received messages on the media to the DLS user until either a `DL_DETACH_REQ` or a `DL_PROMISCOFF_REQ` is received or the stream is closed. Physical level promiscuous mode may be specified for all packets on the medium or for multicast packets only. The stream must be attached to a PPA for these primitives to be accepted.

The `DL_UNITDATA_REQ` primitive is used to send data in a connectionless transfer. Because this is an unacknowledged service, there is no guarantee of delivery. The message consists of one `M_PROTO` message block followed by one or more `M_DATA` blocks containing at least one byte of data.

The `DL_UNITDATA_IND` type is used when a packet is received and is to be passed upstream. The packet is put into an `M_PROTO` message with the primitive set to `DL_UNITDATA_IND`.

The `DL_PHYS_ADDR_REQ` primitive returns the MAC address currently associated with the PPA attached to the stream, in the `DL_PHYS_ADDR_ACK` primitive. When using style 2, this primitive is only valid following a successful `DL_ATTACH_REQ`.

The `DL_SET_PHYS_ADDR_REQ` primitive changes the MAC address currently associated with the PPA attached to the stream. This primitive affects all other current and future streams attached to this device. Once changed, all streams currently or subsequently opened and attached to this device will obtain this new physical address. The new physical address will remain in effect until this primitive is used to change the physical address again or the driver is reloaded.

gld(7D)

Implemented ioctl Functions

The `DL_GET_STATISTICS_REQ` primitive requests a `DL_GET_STATISTICS_ACK` response containing statistics information associated with the PPA attached to the stream. Style 2 streams must be attached to a particular PPA using `DL_ATTACH_REQ` before this primitive will be successful.

GLD supports the `DL_NOTE_LINK_UP`, `DL_NOTE_LINK_DOWN` and `DL_NOTE_SPEED` notifications using the `DL_NOTIFY_IND` primitive. See [dlpi\(7P\)](#).

GLD implements the ioctl `ioc_cmd` function described below. If GLD receives an ioctl command that it does not recognize, it passes it to the device-specific driver's `gldm_ioctl()` routine as described in [gld\(9E\)](#).

The `DLIOCRAW` ioctl function is used by some DLPI applications, most notably the `snoop(1M)` command. The `DLIOCRAW` command puts the stream into a raw mode, which, upon receive, causes the the full MAC-level packet to be sent upstream in an `M_DATA` message instead of it being transformed into the `DL_UNITDATA_IND` form normally used for reporting incoming packets. Packet SAP filtering is still performed on streams that are in raw mode; if a stream user wants to receive all incoming packets it must also select the appropriate promiscuous modes. After successfully selecting raw mode, the application is also allowed to send fully formatted packets to the driver as `M_DATA` messages for transmission. `DLIOCRAW` takes no arguments. Once enabled, the stream remains in this mode until closed.

Requirements on GLD Drivers

GLD-based drivers must include the header file `<sys/gld.h>`.

GLD-based drivers must also specify a link dependency on "misc/gld". (See the `-N` option in [ld\(1\)](#)).

GLD implements the `open(9E)` and `close(9E)` functions and the required `STREAMS` `put(9E)` and `srv(9E)` functions on behalf of the device-specific driver. GLD also implements the `getinfo(9E)` function for the driver.

The `mi_idname` element of the `module_info(9S)` structure is a string specifying the name of the driver. This must exactly match the name of the driver module as it exists in the file system.

The read-side `qinit(9S)` structure should specify the following elements as shown below:

<code>qi_putp</code>	<code>NULL</code>
<code>qi_srvp</code>	<code>gld_rsrv</code>
<code>qi_qopen</code>	<code>gld_open</code>
<code>qi_qclose</code>	<code>gld_close</code>

The write-side `qinit(9S)` structure should specify the following elements as shown below:

<code>qi_putp</code>	<code>gld_wput</code>
----------------------	-----------------------


```

qi_srvp      gld_wsrv
qi_qopen     NULL
qi_qclose    NULL

```

The `devo_getinfo` element of the `dev_ops(9S)` structure should specify `gld_getinfo` as the `getinfo(9E)` routine.

The driver's `attach(9E)` function does all the work of associating the hardware-specific device driver with the GLD facility and preparing the device and driver for use.

The `attach(9E)` function allocates a `gld_mac_info(9S)` ("macinfo") structure using `gld_mac_alloc()`. The driver usually needs to save more information per device than is defined in the macinfo structure; it should allocate the additional required data structure and save a pointer to it in the `gldm_private` member of the `gld_mac_info(9S)` structure.

The `attach(9E)` routine must initialize the macinfo structure as described in `gld_mac_info(9S)` and then call `gld_register()` to link the driver with the GLD module. The driver should map registers if necessary and be fully initialized and prepared to accept interrupts before calling `gld_register()`. The `attach(9E)` function should add interrupts but not enable the device to generate them. The driver should reset the hardware before calling `gld_register()` to ensure it is quiescent; the device must not be started or put into a state where it may generate an interrupt before `gld_register()` is called. That will be done later when GLD calls the driver's `gldm_start()` entry point described in `gld(9E)`. Once `gld_register()` succeeds, the `gld(9E)` entry points may be called by GLD at any time.

The `attach(9E)` routine should return `DDI_SUCCESS` if `gld_register()` succeeds. If `gld_register()` fails, it returns `DDI_FAILURE` and the `attach(9E)` routine should deallocate any resources it allocated before calling `gld_register()` and then also return `DDI_FAILURE`. Under no circumstances should a failed macinfo structure be reused; it should be deallocated using `gld_mac_free()`.

The `detach(9E)` function should attempt to unregister the driver from GLD. This is done by calling `gld_unregister()` described in `gld(9F)`. The `detach(9E)` routine can get a pointer to the needed `gld_mac_info(9S)` structure from the device's private data using `ddi_get_driver_private(9F)`. `gld_unregister()` checks certain conditions that could require that the driver not be detached. If the checks fail, `gld_unregister()` returns `DDI_FAILURE`, in which case the driver's `detach(9E)` routine must leave the device operational and return `DDI_FAILURE`. If the checks succeed, `gld_unregister()` ensures that the device interrupts are stopped, calling the driver's `gldm_stop()` routine if necessary, unlinks the driver from the GLD

gld(7D)

framework, and returns `DDI_SUCCESS`. In this case, the `detach(9E)` routine should remove interrupts, deallocate any data structures allocated in the `attach(9E)` routine, using `gld_mac_free()` to deallocate the `macinfo` structure, and return `DDI_SUCCESS`. It is important to remove the interrupt *before* calling `gld_mac_free()`.

Network Statistics

Solaris network drivers must implement statistics variables. GLD itself tallies some network statistics, but other statistics must be counted by each GLD-based driver. GLD provides support for GLD-based drivers to report a standard set of network driver statistics. Statistics are reported by GLD using the `kstat(7D)` and `kstat(9S)` mechanism. The `DL_GET_STATISTICS_REQ` DLPI command may also be used to retrieve the current statistics counters. All statistics are maintained as unsigned, and all are 32 bits unless otherwise noted.

GLD maintains and reports the following statistics.

<code>rbytes64</code>	Total bytes successfully received on the interface (64 bits).
<code>rbytes</code>	Total bytes successfully received on the interface.
<code>obytes64</code>	Total bytes requested to be transmitted on the interface (64 bits).
<code>obytes</code>	Total bytes requested to be transmitted on the interface.
<code>ipackets64</code>	Total packets successfully received on the interface (64 bits).
<code>ipackets</code>	Total packets successfully received on the interface.
<code>opackets64</code>	Total packets requested to be transmitted on the interface (64 bits).
<code>opackets</code>	Total packets requested to be transmitted on the interface.
<code>multircv</code>	Multicast packets successfully received, including group and functional addresses (long).
<code>multixmt</code>	Multicast packets requested to be transmitted, including group and functional addresses (long).
<code>brdcstrcv</code>	Broadcast packets successfully received (long).
<code>brdcstxmt</code>	Broadcast packets requested to be transmitted (long).
<code>unknowns</code>	Valid received packets not accepted by any stream (long).
<code>noxmtbuf</code>	Packets discarded on output because transmit buffer was busy, or no buffer could be allocated for transmit (long).

<code>blocked</code>	Times a received packet could not be put up a stream because the queue was flow controlled (long).
<code>xmretry</code>	Times transmit was retried after having been delayed due to lack of resources (long).
<code>promisc</code>	Current "promiscuous" state of the interface (string).

The device dependent driver counts the following statistics, keeping track of them in a private per-instance structure. When GLD is asked to report statistics, it calls the driver's `gldm_get_stats()` entry point, as described in `gld(9E)`, to update the device-specific statistics in the `gld_stats(9S)` structure. GLD then reports the updated statistics using the named statistics variables below.

<code>ifspeed</code>	Current estimated bandwidth of the interface in bits per second (64 bits).
<code>media</code>	Current media type in use by the device (string).
<code>intr</code>	Times interrupt handler was called and claimed the interrupt (long).
<code>norcvbuf</code>	Times a valid incoming packet was known to have been discarded because no buffer could be allocated for receive (long).
<code>ierrors</code>	Total packets received that couldn't be processed because they contained errors (long).
<code>oerrors</code>	Total packets that weren't successfully transmitted because of errors (long).
<code>missed</code>	Packets known to have been dropped by the hardware on receive (long).
<code>uflo</code>	Times FIFO underflowed on transmit (long).
<code>oflo</code>	Times receiver overflowed during receive (long).

The following group of statistics applies to networks of type `DL_ETHER`; these are maintained by device-specific drivers of that type, as above.

<code>align_errors</code>	Packets received with framing errors (not an integral number of octets) (long).
<code>fcs_errors</code>	Packets received with CRC errors (long).
<code>duplex</code>	Current duplex mode of the interface (string).
<code>carrier_errors</code>	Times carrier was lost or never detected on a transmission attempt (long).
<code>collisions</code>	Ethernet collisions during transmit (long).

gld(7D)

<code>ex_collisions</code>	Frames where excess collisions occurred on transmit, causing transmit failure (long).
<code>tx_late_collisions</code>	Times a transmit collision occurred late (after 512 bit times) (long).
<code>defer_xmts</code>	Packets without collisions where first transmit attempt was delayed because the medium was busy (long).
<code>first_collisions</code>	Packets successfully transmitted with exactly one collision.
<code>multi_collisions</code>	Packets successfully transmitted with multiple collisions.
<code>sqe_errors</code>	Times SQE test error was reported.
<code>macxmt_errors</code>	Packets encountering transmit MAC failures, except carrier and collision failures.
<code>macrcv_errors</code>	Packets received with MAC errors, except align, fcs, and toolong errors.
<code>toolong_errors</code>	Packets received larger than the maximum permitted length.
<code>runt_errors</code>	Packets received smaller than the minimum permitted length (long).
The following group of statistics applies to networks of type <code>DL_TPR</code> ; these are maintained by device-specific drivers of that type, as above.	
<code>line_errors</code>	Packets received with non-data bits or FCS errors.
<code>burst_errors</code>	Times an absence of transitions for five half-bit timers was detected.
<code>signal_losses</code>	Times loss of signal condition on the ring was detected.
<code>ace_errors</code>	Times an AMP or SMP frame in which A is equal to C is equal to 0, was followed by another such SMP frame without an intervening AMP frame.
<code>internal_errors</code>	Times the station recognized an internal error.
<code>lost_frame_errors</code>	Times the TRR timer expired during transmit.
<code>frame_copied_errors</code>	Times a frame addressed to this station was received with the FS field A bit set to 1.
<code>token_errors</code>	Times the station acting as the active monitor recognized an error condition that needed a token transmitted.
<code>freq_errors</code>	Times the frequency of the incoming signal differed from the expected frequency.

The following group of statistics applies to networks of type `DL_FDDI`; these are maintained by device-specific drivers of that type, as above.

<code>mac_errors</code>	Frames detected in error by this MAC that had not been detected in error by another MAC.
<code>mac_lost_errors</code>	Frames received with format errors such that the frame was stripped.
<code>mac_tokens</code>	Number of tokens received (total of non-restricted and restricted).
<code>mac_tvx_expired</code>	Number of times that TVX has expired.
<code>mac_late</code>	Number of TRT expirations since this MAC was reset or a token was received.
<code>mac_ring_ops</code>	Number of times the ring has entered the "Ring_Operational" state from the "Ring Not Operational" state.

FILES /kernel/misc/gld loadable kernel module

SEE ALSO `ld(1)`, `kstat(7D)`, `dlpi(7P)`, `attach(9E)`, `gld(9E)`, `open(9E)`, `gld(9F)`, `gld_mac_info(9S)`, `gld_stats(9S)`, `kstat(9S)`

Writing Device Drivers

WARNINGS Contrary to the DLPI specification, GLD returns the device's correct address length and broadcast address in `DL_INFO_ACK` even before the stream has been attached to a PPA.

Promiscuous mode may only be entered by streams that are attached to a PPA.

The physical address of a PPA may be changed by the superuser while other streams are bound to the same PPA.

glm(7D)

NAME	glm – GLM SCSI Host Bus Adapter Driver
SYNOPSIS	<code>scsi@unit-address</code>
DESCRIPTION	<p>The <code>glm</code> Host Bus Adapter driver is a SCSI compliant nexus driver that supports the LSI 53c810, LSI 53c875, LSI 53c876, LSI 53C896 and LSI 53C1010 SCSI chips</p> <p>It supports the standard functions provided by the SCSI interface. That is, it supports tagged and untagged queuing, Narrow/Wide/Fast/Ultra SCSI/Ultra SCSI 2/Ultra SCSI 3, and auto request sense, but it does not support linked commands.</p>
Driver Configuration	<p>Configure the <code>glm</code> driver by defining properties in <code>glm.conf</code>. These properties override the global SCSI settings. <code>glm</code> supports these properties which can be modified by the user: <code>scsi-options</code>, <code>target<n>-scsi-options</code>, <code>scsi-reset-delay</code>, <code>scsi-tag-age-limit</code>, <code>scsi-watchdog-tick</code>, and <code>scsi-initiator-id</code>.</p> <p><code>target<n>-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target<n></code>. <code><n></code> can vary from decimal 0 to 15. <code>glm</code> supports these <code>scsi-options</code>: <code>SCSI_OPTIONS_DR</code>, <code>SCSI_OPTIONS_SYNC</code>, <code>SCSI_OPTIONS_TAG</code>, <code>SCSI_OPTIONS_FAST</code>, <code>SCSI_OPTIONS_WIDE</code>, <code>SCSI_OPTIONS_FAST20</code>, <code>SCSI_OPTIONS_FAST40</code> and <code>SCSI_OPTIONS_FAST80</code>.</p> <p>After periodic interval <code>scsi-watchdog-tick</code>, the <code>glm</code> driver searches through all current and disconnected commands for timeouts.</p> <p><code>scsi-tag-age-limit</code> is the number of times that the <code>glm</code> driver attempts to allocate a particular tag ID that is currently in use after going through all tag IDs in a circular fashion. After finding the same tag ID in use <code>scsi-tag-age-limit</code> times, no more commands will be submitted to this target until all outstanding commands complete or timeout.</p> <p>Refer to <code>scsi_hba_attach(9F)</code>.</p>
EXAMPLES	<p>EXAMPLE 1 Using the <code>glm</code> Configuration File</p> <p>Create a file called <code>/kernel/drv/glm.conf</code> and add the following line:</p> <pre>scsi-options=0x78;</pre> <p>This disables tagged queuing, Fast/Ultra SCSI and wide mode for all <code>glm</code> instances.</p> <p>The following example disables an option for one specific <code>glm</code> (refer to <code>driver.conf(4)</code> and <code>pci(4)</code> for more details):</p> <pre>name="glm" parent="/pci@1f,4000" unit-address="3" target1-scsi-options=0x58 scsi-options=0x178 scsi-initiator-id=6;</pre> <p>Note that the default initiator ID in OBP is 7 and that the change to ID 6 will occur at attach time. It may be preferable to change the initiator ID in OBP.</p>

EXAMPLE 1 Using the glm Configuration File (Continued)

The example above sets `scsi-options` for target 1 to 0x58 and all other targets on this SCSI bus to 0x178.

The physical pathname of the parent can be determined using the `/devices` tree or following the link of the logical device name:

```
# ls -l /dev/rdisk/c0t0d0s0
lrwxrwxrwx 1 root root 45 May 16 10:08 /dev/rdisk/c0t0d0s0 ->
  . . / . . /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

In this case, like the example above, the parent is `/pci@1f,4000` and the `unit-address` is the number bound to the `scsi@3` node.

To set `scsi-options` more specifically per target:

```
target1-scsi-options=0x78;
device-type-scsi-options-list =
  "SEAGATE ST32550W", "seagate-scsi-options" ;
seagate-scsi-options = 0x58;
scsi-options=0x3f8;
```

The above sets `scsi-options` for target 1 to 0x78 and for all other targets on this SCSI bus to 0x3f8 except for one specific disk type which will have `scsi-options` set to 0x58.

`scsi-options` specified per target ID have the highest precedence, followed by `scsi-options` per device type. Global `scsi-options` (for all glm instances) per bus have the lowest precedence.

The system needs to be rebooted before the specified `scsi-options` take effect.

Driver Capabilities

The target driver needs to set capabilities in the glm driver in order to enable some driver features. The target driver can query and modify these capabilities: `synchronous`, `tagged-qing`, `wide-xfer`, `auto-rqsense`, `qfull-retries`, `qfull-retry-interval`. All other capabilities can only be queried.

By default, `tagged-qing`, `auto-rqsense`, and `wide-xfer` capabilities are disabled, while `disconnect`, `synchronous`, and `untagged-qing` are enabled. These capabilities can only have binary values (0 or 1). The default value for `qfull-retries` is 10 and the default value for `qfull-retry-interval` is 100. The `qfull-retries` capability is a `uchar_t` (0 to 255) while `qfull-retry-interval` is a `ushort_t` (0 to 65535).

The target driver needs to enable `tagged-qing` and `wide-xfer` explicitly. The `untagged-qing` capability is always enabled and its value cannot be modified.

Whenever there is a conflict between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only whom `!= 0` is supported in the `scsi_ifsetcap(9F)` call.

glm(7D)

Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

FILES /kernel/drv/glm ELF Kernel Module
/kernel/drv/glm.conf Optional configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Limited to PCI-based systems with LSI 53c810, LSI 53c875, LSI 53c876, LSI 53c896 and LSI 53c1010 SCSI I/O processors

SEE ALSO `prtconf(1M)`, `driver.conf(4)`, `pci(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

ANSI Small Computer System Interface-2 (SCSI-2),

LSI Logi Inc (formerly Symbios Logic Inc.):

- SYM53c810 PCI-SCSI I/O processor with Narrow operation
- SYM53c875 PCI-SCSI I/O Processor With Fast-20
- SYM53c876 PCI-SCSI I/O processor Dual channel Fast-20
- SYM53c896 PCI-SCSI I/O processor Dual channel Fast-40
- SYM53c1010 PCI-SCSI I/O processor Dual Channel Fast-80

DIAGNOSTICS The messages described below are some that may appear on the system console, as well as being logged.

Device is using a hilevel intr

The device was configured with an interrupt level that cannot be used with this `glm` driver. Check the PCI device.

map setup failed

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

glm_script_alloc failed

The driver was unable to load the SCRIPTS for the SCSI processor, check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

cannot map configuration space.

The driver was unable to map in the configuration registers. Check for bad hardware. SCSI devices will be inaccessible.

attach failed

The driver was unable to attach; usually preceded by another warning that indicates why attach failed. These can be considered hardware failures.

SCSI bus DATA IN phase parity error	The driver detected parity errors on the SCSI bus.
SCSI bus MESSAGE IN phase parity error	The driver detected parity errors on the SCSI bus.
SCSI bus STATUS phase parity error	The driver detected parity errors on the SCSI bus.
Unexpected bus free	Target disconnected from the bus without notice. Check for bad hardware.
Disconnected command timeout for Target <id>.<lun>	A timeout occurred while target <i>id/lun</i> was disconnected. This is usually a target firmware problem. For tagged queuing targets, <n> commands were outstanding when the timeout was detected.
Disconnected tagged cmd(s) (<n>) timeout for Target <id>.<lun>	A timeout occurred while target <i>id/lun</i> was disconnected. This is usually a target firmware problem. For tagged queuing targets, <n> commands were outstanding when the timeout was detected.
Connected command timeout for Target <id>.<lun>	This is usually a SCSI bus problem. Check cables and termination.
Target <id> reducing sync. transfer rate	A data transfer hang or DATA-IN phase parity error was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.
Target <id> reverting to async. mode	A second data transfer hang was detected for this target. The driver attempts to eliminate this problem by reducing the data transfer rate.
Target <id> disabled wide SCSI mode	A second data phase hang was detected for this target. The driver attempts to eliminate this problem by disabling wide SCSI mode.
auto request sense failed	An attempt to start an auto request packet failed. Another auto request packet may already be in transport.
invalid reselection (<id>.<lun>)	A reselection failed; target accepted <code>abort</code> or <code>reset</code> , but still tries to reconnect. Check for bad hardware.
invalid intcode	The SCRIPTS processor generated an invalid SCRIPTS interrupt. Check for bad hardware.

NOTES The glm driver supports the following LSI chips:

- LSI 53C810, which supports Narrow, Fast SCSI mode. The maximum SCSI bandwidth is 10 MB/sec.
- LSI 53C875, which supports Wide, Fast, and Ultra SCSI mode. The maximum SCSI bandwidth is 40 MB/sec.

glm(7D)

- LSI 53C896, which supports Wide, Fast and Ultra SCSI 2 mode. The maximum LVD SCSI bandwidth is 80 MB/sec.
- LSI 53c1010, which supports wide, Fast and Ultra SCSI 3 mode. The maximum LVD SCSI bandwidth is 160 MB/sec.

The `glm` driver exports properties indicating per target the negotiated transfer speed (`target<n>-sync-speed`), whether wide bus is supported (`target<n>-wide`), for that particular target (`target<n>-scsi-options`), and whether tagged queuing has been enabled (`target<n>-TQ`). The `sync-speed` property value is the data transfer rate in KB/sec. The `target<n>-TQ` and the `target<n>-wide` property have value 1 to indicate that the corresponding capability is enabled, or 0 to indicate that the capability is disabled for that target. Refer to `prtconf(1M)` (verbose option) for viewing the `glm` properties.

```
scsi, instance #0
  Driver properties:
    name <target6-TQ> length <4>
      value <0x00000000>.
    name <target6-wide> length <4>
      value <0x00000000>.
    name <target6-sync-speed> length <4>
      value <0x00002710>.
    name <target1-TQ> length <4>
      value <0x00000001>.
    name <target1-wide> length <4>
      value <0x00000000>.
    name <target1-sync-speed> length <4>
      value <0x00002710>.
    name <target0-TQ> length <4>
      value <0x00000001>.
    name <target0-wide> length <4>
      value <0x00000001>.
    name <target0-sync-speed> length <4>
      value <0x00009c40>.
    name <scsi-options> length <4>
      value <0x000007f8>.
    name <scsi-watchdog-tick> length <4>
      value <0x0000000a>.
    name <scsi-tag-age-limit> length <4>
      value <0x00000002>.
    name <scsi-reset-delay> length <4>
      value <0x00000bb8>.
    name <latency-timer> length <4>
      value <0x00000088>.
    name <cache-line-size> length <4>
      value <0x00000010>.
```

gpio_87317(7D)

NAME | gpio_87317 – General purpose I/O driver for SuperIO

DESCRIPTION | The `gpio_87317` driver is the general purpose I/O driver for the National Semiconductor SuperIO (PC87317) chipset. It supports remote system controller (RSC) administration via an interface to the SuperIO's general purpose I/O bits.

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Limited to SPARC systems with SuperIO
Availability	SUNWcar
Interface Stability	Unstable

SEE ALSO | *PC87317VUL/PC97317VUL SuperI/O Data Sheet* — National Semiconductor

grbeep(7d)

- NAME** | grbeep – Platform-dependent beep driver for SMBus-based hardware
- SYNOPSIS** | beep@unit-address
- DESCRIPTION** | The grbeep driver generates beeps on platforms (including Sun Blade 100, 150, 1500, 2500) that use SMBbus-based registers and USB keyboards. When the `KIOCCMD` ioctl is issued to the USB keyboard module (see `usbkbm(7M)`) with command `KBD_CMD_BELL/KBD_CMD_NOBELL`, `usbkbm(7M)` passes the request to the grbeep driver to turn the beep on and off, respectively.
- FILES** | /platform/sun4u/kernel/drv/sparcv9/grbeep
64-bit ELF kernel driver
- ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SMBus-based SPARC
Availability	SUNWcarx.u

- SEE ALSO** | `kbd(1)`, `attributes(5)`, `bbc_beep(7D)`, `kb(7M)`, `usbkbm(7M)`
Writing Device Drivers
- DIAGNOSTICS** | None

NAME	hci1394 – 1394 OpenHCI host controller driver								
SYNOPSIS	firewire@unit-address								
DESCRIPTION	The hci1394 host controller driver is an IEEE 1394 compliant nexus driver that supports the <i>1394 Open Host Controller Interface Specification 1.0</i> , an industry standard developed by Sun, Apple, Compaq, Intel, Microsoft, National Semiconductor, and Texas Instruments. The hci1394 driver supports asynchronous transfers, isochronous transfers, and bus reset management. The hci1394 driver also supports the nexus device control interface.								
FILES	<p>/kernel/drv/sparcv9/hci1394 64-bit SPARC ELF kernel module</p> <p>/kernel/drv/hci1394 32-bit x86 ELF kernel module</p> <p>/kernel/drv/amd64/hci1394 64-bit x86 ELF kernel module</p>								
ATTRIBUTES	See attributes(5) for a description of the following attributes:								
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>SPARC, x86, PCI-based systems</td> </tr> <tr> <td>Availability</td> <td>SUNW1394x</td> </tr> <tr> <td>Interface Stability</td> <td>Unstable</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC, x86, PCI-based systems	Availability	SUNW1394x	Interface Stability	Unstable
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Architecture	SPARC, x86, PCI-based systems								
Availability	SUNW1394x								
Interface Stability	Unstable								
SEE ALSO	<p>IEEE 1394 - IEEE Standard for a High Performance Serial Bus</p> <p>1394 Open Host Controller Interface Specification 1.0</p>								

hdio(7I)

NAME	hdio – SMD and IPI disk control operations																
SYNOPSIS	<pre>#include <sys/hdio.h></pre>																
DESCRIPTION	<p>Note – The SMC and IPI drivers have been discontinued. dkio(7I) is now the preferred method for retrieving disk information.</p> <p>The SMD and IPI disk drivers supplied with this release support a set of <code>ioctl(2)</code> requests for diagnostics and bad sector information. Basic to these <code>ioctl()</code> requests are the definitions in <code><sys/hdio.h></code>.</p>																
IOCTLS	<table><tr><td>HDKIOCGTYPE</td><td>The argument is a pointer to a <code>hdk_type</code> structure (described below). This <code>ioctl()</code> gets specific information from the hard disk.</td></tr><tr><td>HDKIOCSTYPE</td><td>The argument is a pointer to a <code>hdk_type</code> structure (described below). This <code>ioctl()</code> sets specific information about the hard disk.</td></tr><tr><td colspan="2"><pre>/* * Used for drive info */ struct hdk_type { ushort_t hdk_t_hsect; /* hard sector count (read only) */ ushort_t hdk_t_promrev; /* prom revision (read only) */ uchar_t hdk_t_drtype; /* drive type (ctlr specific) */ uchar_t hdk_t_drstat; /* drive status (ctlr specific, ro) */ };</pre></td></tr><tr><td>HDKIOCGBAD</td><td>The argument is a pointer to a <code>hdk_badmap</code> structure (described below). This <code>ioctl()</code> is used to get the bad sector map from the disk.</td></tr><tr><td>HDKIOCSBAD</td><td>The argument is a pointer to a <code>hdk_badmap</code> structure (described below). This <code>ioctl()</code> is used to set the bad sector map on the disk.</td></tr><tr><td colspan="2"><pre>/* * Used for bad sector map */ struct hdk_badmap { caddr_t hdkb_bufaddr; /* address of user's map buffer */ };</pre></td></tr><tr><td>HDKIOCGDIAG</td><td>The argument is a pointer to a <code>hdk_diag</code> structure (described below). This <code>ioctl()</code> gets the most recent command that failed along with the sector and error number from the hard disk.</td></tr><tr><td colspan="2"><pre>/* * Used for disk diagnostics */ struct hdk_diag { ushort_t hdkd_errcmd; /* most recent command in error */ daddr_t hdkd_errsect; /* most recent sector in error */ uchar_t hdkd_errno; /* most recent error number */ uchar_t hdkd_severe; /* severity of most recent error */ };</pre></td></tr></table>	HDKIOCGTYPE	The argument is a pointer to a <code>hdk_type</code> structure (described below). This <code>ioctl()</code> gets specific information from the hard disk.	HDKIOCSTYPE	The argument is a pointer to a <code>hdk_type</code> structure (described below). This <code>ioctl()</code> sets specific information about the hard disk.	<pre>/* * Used for drive info */ struct hdk_type { ushort_t hdk_t_hsect; /* hard sector count (read only) */ ushort_t hdk_t_promrev; /* prom revision (read only) */ uchar_t hdk_t_drtype; /* drive type (ctlr specific) */ uchar_t hdk_t_drstat; /* drive status (ctlr specific, ro) */ };</pre>		HDKIOCGBAD	The argument is a pointer to a <code>hdk_badmap</code> structure (described below). This <code>ioctl()</code> is used to get the bad sector map from the disk.	HDKIOCSBAD	The argument is a pointer to a <code>hdk_badmap</code> structure (described below). This <code>ioctl()</code> is used to set the bad sector map on the disk.	<pre>/* * Used for bad sector map */ struct hdk_badmap { caddr_t hdkb_bufaddr; /* address of user's map buffer */ };</pre>		HDKIOCGDIAG	The argument is a pointer to a <code>hdk_diag</code> structure (described below). This <code>ioctl()</code> gets the most recent command that failed along with the sector and error number from the hard disk.	<pre>/* * Used for disk diagnostics */ struct hdk_diag { ushort_t hdkd_errcmd; /* most recent command in error */ daddr_t hdkd_errsect; /* most recent sector in error */ uchar_t hdkd_errno; /* most recent error number */ uchar_t hdkd_severe; /* severity of most recent error */ };</pre>	
HDKIOCGTYPE	The argument is a pointer to a <code>hdk_type</code> structure (described below). This <code>ioctl()</code> gets specific information from the hard disk.																
HDKIOCSTYPE	The argument is a pointer to a <code>hdk_type</code> structure (described below). This <code>ioctl()</code> sets specific information about the hard disk.																
<pre>/* * Used for drive info */ struct hdk_type { ushort_t hdk_t_hsect; /* hard sector count (read only) */ ushort_t hdk_t_promrev; /* prom revision (read only) */ uchar_t hdk_t_drtype; /* drive type (ctlr specific) */ uchar_t hdk_t_drstat; /* drive status (ctlr specific, ro) */ };</pre>																	
HDKIOCGBAD	The argument is a pointer to a <code>hdk_badmap</code> structure (described below). This <code>ioctl()</code> is used to get the bad sector map from the disk.																
HDKIOCSBAD	The argument is a pointer to a <code>hdk_badmap</code> structure (described below). This <code>ioctl()</code> is used to set the bad sector map on the disk.																
<pre>/* * Used for bad sector map */ struct hdk_badmap { caddr_t hdkb_bufaddr; /* address of user's map buffer */ };</pre>																	
HDKIOCGDIAG	The argument is a pointer to a <code>hdk_diag</code> structure (described below). This <code>ioctl()</code> gets the most recent command that failed along with the sector and error number from the hard disk.																
<pre>/* * Used for disk diagnostics */ struct hdk_diag { ushort_t hdkd_errcmd; /* most recent command in error */ daddr_t hdkd_errsect; /* most recent sector in error */ uchar_t hdkd_errno; /* most recent error number */ uchar_t hdkd_severe; /* severity of most recent error */ };</pre>																	

hdio(7I)

};

SEE ALSO ioctl(2), dkio(7I)

hid(7D)

NAME	hid – Human interface device (HID) class driver
SYNOPSIS	keyboard@unit-address mouse@unit-address input@unit-address:consumer_control
DESCRIPTION	<p>The hid driver is a USBA (Solaris USB Architecture) compliant client driver that supports the <i>Human Interface Device Class (HID) 1.0</i> specification. The Human Interface Device (HID) class encompasses devices controlled by humans to operate computer systems. Typical examples of HID devices include keyboards, mice, trackballs, and joysticks. HID also covers front-panel controls such as knobs, switches, and buttons. A USB device with multiple interfaces may have one interface for audio and a HID interface to define the buttons that control the audio.</p> <p>The hid driver is general and primarily handles the USB functionality of the device and generic HID functionality. For example, HID interfaces are required to have an interrupt pipe for the device to send data packets, and the hid driver opens the pipe to the interrupt endpoint and starts polling. The hid driver is also responsible for managing the device through the default control pipe. In addition to being a USB client driver, the hid driver is also a STREAMS driver so that modules may be pushed on top of it.</p> <p>The HID specification is flexible, and HID devices dynamically describe their packets and other parameters through a HID report descriptor. The HID parser is a misc module that parses the HID report descriptor and creates a database of information about the device. The hid driver queries the HID parser to find out the type and characteristics of the HID device. The HID specification predefines packet formats for the boot protocol keyboard and mouse.</p>
FILES	/kernel/drv/hid 32-bit x86 ELF kernel hid module /kernel/drv/amd64/hid 64-bit x86 ELF kernel hid module /kernel/drv/sparcv9/hid 64-bit SPARC ELF kernel hid module /kernel/misc/hidparser 32-bit x86 ELF kernel hidparser module /kernel/misc/amd64/hidparser 64-bit x86 ELF kernel hidparser module /kernel/misc/sparcv9/hidparser 64-bit SPARC ELF kernel hidparser module
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

hid(7D)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86, PCI-based systems
Availability	SUNWusb

SEE ALSO [cfgadm_usb\(1M\)](#), [attributes\(5\)](#), [usba\(7D\)](#)

Writing Device Drivers

STREAMS Programming Guide

Universal Serial Bus Specification 1.0 and 1.1

Device Class Definition for Human Interface Devices (HID) 1.1

System Administration Guide: Basic Administration

<http://www.sun.com/io>

DIAGNOSTICS None.

NOTES The hid driver currently supports only keyboard, mouse and audio HID control device.

hme(7D)

NAME	hme – SUNW,hme Fast-Ethernet device driver
SYNOPSIS	/dev/hme
DESCRIPTION	<p>The SUNW, hme Fast-Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, dlpi(7P), over a SUNW, hme Fast-Ethernet controller. The motherboard and add-in SBus SUNW, hme controllers of several varieties are supported. Multiple SUNW, hme controllers installed within the system are supported by the driver.</p> <p>The hme driver provides basic support for the SUNW, hme hardware. It is used to handle the SUNW, hme device. Functions include chip initialization, frame transit and receive, multicast and promiscuous support, and error recovery and reporting.</p> <p>SUNW, hme The SUNW, hme device provides 100Base-TX networking interfaces using SUN's FEPS ASIC and an Internal Transceiver. The FEPS ASIC provides the Sbus interface and MAC functions and the Physical layer functions are provided by the Internal Transceiver which connects to a RJ-45 connector. In addition to the RJ-45 connector, an MII (Media Independent Interface) connector is also provided on all SUNW, hme devices except the SunSwith SBus adapter board. The MII interface is used to connect to an External Transceiver which may use any physical media (copper or fiber) specified in the 100Base-TX standard. When an External Transceiver is connected to the MII, the driver selects the External Transceiver and disables the Internal Transceiver.</p> <p>The 100Base-TX standard specifies an "auto-negotiation" protocol to automatically select the mode and speed of operation. The Internal transceiver is capable of doing "auto-negotiation" with the remote-end of the link (Link Partner) and receives the capabilities of the remote end. It selects the Highest Common Denominator mode of operation based on the priorities. It also supports forced-mode of operation where the driver can select the mode of operation.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device /dev/hme is used to access all SUNW, hme controllers installed within the system.</p> <p>The hme driver is a "style 2" Data Link Service provider. All M_PROTO and M_PCPROTO type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <sys/dlpi.h>. Refer to dlpi(7P) for more information. An explicit DL_ATTACH_REQ message by the user is required to associate the opened stream with a particular device (ppa). The ppa ID is interpreted as an unsigned long data type and indicates the corresponding device instance (unit) number. An error (DL_ERROR_ACK) is returned by the driver if the ppa field value does not correspond to a valid device instance number for this system. The device is initialized on first attach and de-initialized (stopped) at last detach.</p> <p>The values returned by the driver in the DL_INFO_ACK primitive in response to the DL_INFO_REQ from the user are as follows:</p> <ul style="list-style-type: none">■ The maximum SDU is 1500 (ETHERMTU - defined in <sys/ethernet.h>).■ The minimum SDU is 0.■ The dl_sap address length is 8.

- The MAC type is `DL_ETHER`.
- The `sap` length values is `-2` meaning the physical address component is followed immediately by a 2 byte `sap` component within the DLSAP address.
- The service mode is `DL_CLDLS`.
- No optional quality of service (QOS) support is included at present so the QOS fields are 0.
- The provider style is `DL_STYLE2`.
- The version is `DL_VERSION_2`.
- The broadcast address value is Ethernet/IEEE broadcast address (`0xFFFFFFFF`).

Once in the `DL_ATTACHED` state, the user must send a `DL_BIND_REQ` to associate a particular SAP (Service Access Pointer) with the stream. The hme driver interprets the `sap` field within the `DL_BIND_REQ` as an Ethernet “type” therefore valid values for the `sap` field are in the `[0-0xFFFFFFFF]` range. Only one Ethernet type can be bound to the stream at any time.

If the user selects a `sap` with a value of 0, the receiver will be in “802.3 mode”. All frames received from the media having a “type” field in the range `[0-1500]` are assumed to be 802.3 frames and are routed up all open Streams which are bound to `sap` value 0. If more than one Stream is in “802.3 mode” then the frame will be duplicated and routed up multiple Streams as `DL_UNITDATA_IND` messages.

In transmission, the driver checks the `sap` field of the `DL_BIND_REQ` if the `sap` value is 0, and if the destination type field is in the range `[0-1500]`. If either is true, the driver computes the length of the message, not including initial `M_PROTO` mblk (message block), of all subsequent `DL_UNITDATA_REQ` messages and transmits 802.3 frames that have this value in the MAC frame header length field.

The hme driver DLSAP address format consists of the 6 byte physical (Ethernet) address component followed immediately by the 2 byte `sap` (type) component producing an 8 byte DLSAP address. Applications should *not* hardcode to this particular implementation-specific DLSAP address format but use information returned in the `DL_INFO_ACK` primitive to compose and decompose DLSAP addresses. The `sap` length, full DLSAP length, and `sap/physical` ordering are included within the `DL_INFO_ACK`. The physical address length can be computed by subtracting the `sap` length from the full DLSAP address length or by issuing the `DL_PHYS_ADDR_REQ` to obtain the current physical address associated with the stream.

Once in the `DL_BOUND` state, the user may transmit frames on the Ethernet by sending `DL_UNITDATA_REQ` messages to the hme driver. The hme driver will route received Ethernet frames up all those open and bound streams having a `sap` which matches the Ethernet type as `DL_UNITDATA_IND` messages. Received Ethernet frames are duplicated and routed up multiple open streams if necessary. The DLSAP address contained within the `DL_UNITDATA_REQ` and `DL_UNITDATA_IND` messages consists of both the `sap` (type) and physical (Ethernet) components.

hme(7D)

hme Primitives

In addition to the mandatory connectionless DLPI message set the driver additionally supports the following primitives.

The `DL_ENABMULTI_REQ` and `DL_DISABMULTI_REQ` primitives enable/disable reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. These primitives are accepted by the driver in any state following `DL_ATTACHED`.

The `DL_PROMISCON_REQ` and `DL_PROMISCOFF_REQ` primitives with the `DL_PROMISC_PHYS` flag set in the `d1_level` field enables/disables reception of all (“promiscuous mode”) frames on the media including frames generated by the local host. When used with the `DL_PROMISC_SAP` flag set this enables/disables reception of all `sap` (Ethernet type) values. When used with the `DL_PROMISC_MULTI` flag set this enables/disables reception of all multicast group addresses. The effect of each is always on a per-stream basis and independent of the other `sap` and physical level configurations on this stream or other streams.

The `DL_PHYS_ADDR_REQ` primitive returns the 6 octet Ethernet address currently associated (attached) to the stream in the `DL_PHYS_ADDR_ACK` primitive. This primitive is valid only in states following a successful `DL_ATTACH_REQ`.

The `DL_SET_PHYS_ADDR_REQ` primitive changes the 6 octet Ethernet address currently associated (attached) to this stream. The credentials of the process which originally opened this stream must be superuser. Otherwise `EPERM` is returned in the `DL_ERROR_ACK`. This primitive is destructive in that it affects all other current and future streams attached to this device. An `M_ERROR` is sent up all other streams attached to this device when this primitive is successful on this stream. Once changed, all streams subsequently opened and attached to this device will obtain this new physical address. Once changed, the physical address will remain until this primitive is used to change the physical address again or the system is rebooted, whichever comes first.

hme DRIVER

By default, the hme driver performs “auto-negotiation” to select the mode and speed of the link, when the Internal Transceiver is used.

When an External Transceiver is connected to the MII interface, the driver selects the External Transceiver for networking operations. If the External Transceiver supports “auto-negotiation”, the driver uses the auto-negotiation procedure to select the link speed and mode. If the External Transceiver does not support auto-negotiation, it will select the highest priority mode supported by the transceiver.

- 100 Mbps, full-duplex
- 100 Mbps, half-duplex
- 10 Mbps, full-duplex
- 10 Mbps, half-duplex

The link can be in one of the 4 following modes:

These speeds and modes are described in the 100Base-TX standard.

The *auto-negotiation* protocol automatically selects:

- Operation mode (half-duplex or full-duplex)
- Speed (100 Mbps or 10 Mbps)

The auto-negotiation protocol does the following:

- Gets all the modes of operation supported by the Link Partner
- Advertises its capabilities to the Link Partner
- Selects the highest common denominator mode of operation based on the priorities

The *internal transceiver* is capable of all of the operating speeds and modes listed above. When the internal transceiver is used, by *default*, auto-negotiation is used to select the speed and the mode of the link and the common mode of operation with the Link Partner.

When an *external transceiver* is connected to the MII interface, the driver selects the external transceiver for networking operations. If the external transceiver supports auto-negotiation:

- The driver uses the auto-negotiation procedure to select the link speed and mode.

If the external transceiver *does not* support auto-negotiation

- The driver selects the highest priority mode supported by the transceiver.

Sometimes, the user may want to select the speed and mode of the link. The SUNW, hme device supports programmable "IPG" (Inter-Packet Gap) parameters `ipg1` and `ipg2`. By default, the driver sets `ipg1` to 8 `byte-times` and `ipg2` to 4 `byte-times` (which are the standard values). Sometimes, the user may want to alter these values depending on whether the driver supports 10 Mbps or 100 Mbps and accordingly, IPG will be set to 9.6 or 0.96 microseconds.

hme Parameter List

The hme driver provides for setting and getting various parameters for the SUNW, hme device. The parameter list includes:

```
current transceiver status
current link status
inter-packet gap
local transceiver capabilities
link partner capabilities
```

The local transceiver has two set of capabilities: one set reflects the capabilities of the hardware, which are `read-only` (RO) parameters and the second set reflects the values chosen by the user and is used in speed selection. There are `read/write` (RW) capabilities. At boot time, these two sets of capabilities will be the same. The Link Partner capabilities are also read only parameters because the current default value of these parameters can only be read and cannot be modified.

FILES	<code>/dev/hme</code>	hme special character device
	<code>/kernel/drv/hme.conf</code>	System-wide default device driver properties

hme(7D)

SEE ALSO ndd(1M), netstat(1M), driver.conf(4), dlpi(7P)

NAME hpfc – Agilent fibre channel host bus adapter

SYNOPSIS PCI pci103c

DESCRIPTION The hpfc fibre channel host bus adapter is a SCSI compliant nexus driver that supports all Agilent fibre channel host bus adapters, including the HHBA5100x, HHBA5101x, and HHBA5121x models. Agilent host bus adapters support the fibre channel protocol on private fibre channel arbitrated loops and fabrics. The driver supports up to ten host bus adapters, with a maximum of 125 fibre channel devices on each host bus adapter. The hpfc driver supports a maximum of 256 LUNs per target.

The hpfc driver does not support the BIOS Int 13 feature, which enables the booting of an operating system. As a result, you should not install an operating system on devices attached to the hpfc driver.

CONFIGURATION The hpfc driver attempts to configure itself using the information in the /kernel/drv/hpfc.conf configuration file.

By default, the driver supports only LUN 0 for each target device. To add multiple LUN support, modify the /kernel/drv/sd.conf file.

Before upgrading the hpfc driver, backup the sd.conf file to save customized LUN settings and then use pkgrm(1M) to remove the old version of the driver.

The host bus adapter port is initialized to FL_Port when connected to a fabric switch. To change it to F_Port, add the init_as_nport=1 entry to the hpfc.conf file and reboot the system.

To conserve system resources, at least one disk drive must be attached to the hpfc driver. If no devices are attached, the driver will not load.

FILES

/kernel/drv/hpfc	32-bit ELF kernel module
/kernel/drv/sparcv9/hpfc	64-bit ELF kernel module
/kernel/drv/hpfc.conf	Driver configuration file
/kernel/drv/sd.conf	SCSI disk configuration file

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86, SPARC

SEE ALSO luxadm(1M), pkgrm(1M), prtconf(1M), driver.conf(4), attributes(5), ses(7D), ssd(7D)

ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL),

ANSI X3.269-1996, Fibre Channel Protocol for SCSI (FCP),

hpfc(7D)

*ANSI X3.270-1996, SCSI-3 Architecture Model (SAM),
Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)*

NAME	hsfs – High Sierra & ISO 9660 CD-ROM file system
DESCRIPTION	<p>HSFS is a file system type that allows users access to files on High Sierra or ISO 9660 format CD-ROM disks from within the SunOS operating system. Once mounted, a HSFS file system provides standard SunOS read-only file system operations and semantics. That is, users can read files and list files in a directory on a High Sierra or ISO 9660 CD-ROM, and applications can use standard UNIX system calls on these files and directories.</p> <p>This file system also contains support for the Rock Ridge Extensions. If the extensions are contained on the CD-ROM, then the file system will provide all of the file system semantics and file types of UFS, except for writability and hard links.</p> <p>If your <code>/etc/vfstab</code> file contains a line similar to</p> <pre style="margin-left: 20px;">/dev/dsk/c0t6d0s0 -/hsfs hsfs -no roand /hsfs exists, you can mount an HSFS file system with either of the following commands:</pre> <pre style="margin-left: 20px;">mount -F hsfs -o ro device-special directory-nameOR</pre> <pre style="margin-left: 20px;">mount /hsfs</pre> <p>Normally, if Rock Ridge extensions exist on the CD-ROM, the file system will automatically use those extensions. If you do not want to use the Rock Ridge extensions, use the “<code>nrr</code>” (No Rock Ridge) mount option. The mount command would then be:</p> <pre style="margin-left: 20px;">mount -F hsfs -o ro,nrr device-special directory-name</pre> <p>Files on a High Sierra or ISO 9660 CD-ROM disk have names of the form <code>filename.ext;version</code>, where <code>filename</code> and the optional <code>ext</code> consist of a sequence of uppercase alphanumeric characters (including “<code>_</code>”), while the <code>version</code> consists of a sequence of digits, representing the version number of the file. HSFS converts all the uppercase characters in a file name to lowercase, and truncates the “<code>;</code>” and version information. If more than one version of a file is present on the CD-ROM, only the file with the highest version number is accessible.</p> <p>Conversion of uppercase to lowercase characters may be disabled by using the <code>-o nomap1case</code> option to <code>mount(1M)</code>. (See <code>mount_hsfs(1M)</code>).</p> <p>If the CD-ROM contains Rock Ridge extensions, the file names and directory names may contain any character supported under UFS. The names may also be upper and/or lower case and will be case sensitive. File name lengths can be as long as those of UFS.</p> <p>Files accessed through HSFS have mode 555 (owner, group and world readable and executable), uid 0 and gid 3. If a directory on the CD-ROM has read permission, HSFS grants execute permission to the directory, allowing it to be searched.</p> <p>With Rock Ridge extensions, files and directories can have any permissions that are supported on a UFS file system; however, despite any write permissions, the file system is read-only, with <code>EROFS</code> returned to any write operations.</p>

hsfs(7FS)

High Sierra and ISO 9660 CD-ROMs support only regular files and directories, thus HSFS supports only these file types. A Rock Ridge CD-ROM can support regular files, directories, and symbolic links, as well as device nodes, such as block, character, and FIFO.

EXAMPLES **EXAMPLE 1** Sample Display of File System Files

If there is a file `BIG.BAR` on a High Sierra or ISO 9660 format CD-ROM it will show up as `big.bar` when listed on a HSFS file system.

If there are three files

```
BAR.BAZ;1
```

```
BAR.BAZ;2
```

and

```
BAR.BAZ;3
```

on a High Sierra or ISO 9660 format CD-ROM, only the file `BAR.BAZ;3` will be accessible. It will be listed as `bar.baz`.

SEE ALSO `mount(1M)`, `mount_hfs(1M)`, `vfstab(4)`

N. V. Phillips and Sony Corporation, *System Description Compact Disc Digital Audio*, ("Red Book").

N. V. Phillips and Sony Corporation, *System Description of Compact Disc Read Only Memory*, ("Yellow Book").

IR "Volume and File Structure of CD-ROM for Information Interchange", ISO 9660:1988(E).

DIAGNOSTICS

hsfs: Warning: the file system...

does not conform to the ISO-9660 spec

The specific reason appears on the following line. You might be attempting to mount a CD-ROM containing a different file system, such as UFS.

hsfs: Warning: the file system...

contains a file [with an] unsupported type

The `hsfs` file system does not support the format of some file or directory on the CD-ROM, for example a record structured file.

hsfs: hsnode table full, %d nodes allocated

There are not enough HSFS internal data structure elements to handle all the files currently open. This problem may be overcome by adding a line of the form `set hsfs:nhnode=number` to the `/etc/system` system configuration file and rebooting. See `system(4)`.

WARNINGS

Do not physically eject a CD-ROM while the device is still mounted as a HSFS file system.

Under MS-DOS (for which CD-ROMs are frequently targeted), files with no extension may be represented either as

filename.Or

filename that is, with or without a trailing period. These names are not equivalent under UNIX systems. For example, the names

BAR. and

BAR are not names for the same file under the UNIX system. This may cause confusion if you are consulting documentation for CD-ROMs originally intended for MS-DOS systems.

Use of the `-o notraildot` option to `mount(1M)` makes it optional to specify the trailing dot. (See `mount_hfs(1M)`).

NOTES No translation of any sort is done on the contents of High Sierra or ISO 9660 format CD-ROMs; only directory and file names are subject to interpretation by HSFS.

hubd(7D)

NAME hubd – USB hub driver

SYNOPSIS hub@unit-address

DESCRIPTION The hubd is a USB A (Solaris USB Architecture) compliant client driver that supports USB hubs conforming to the *Universal Serial Bus Specification 2.0*. The hubd driver supports bus-powered and self-powered hubs. The driver supports hubs with individual port power, ganged power and no power switching.

When a device is attached to a hub port, the hubd driver enumerates the device by determining its type and assigning an address to it. For multi-configuration devices, hubd sets the preferred configuration (refer to `cfgadm_usb(1M)` to select a configuration). The hubd driver attaches a driver to the device if one is available for the default or selected configuration. When the device is disconnected from the hub port, the hubd driver offlines any driver instance attached to the device.

FILES

- `/kernel/drv/hubd`
32-bit x86 ELF kernel module
- `/kernel/drv/amd64/hubd`
64-bit x86 ELF kernel module
- `/kernel/drv/sparcv9/hubd`
64-bit SPARC ELF kernel module

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86, PCI-based systems
Availability	SUNWusb

SEE ALSO `cfgadm_usb(1M)`, `attributes(5)`, `usba(7D)`

Writing Device Drivers

Universal Serial Bus Specification 2.0

System Administration Guide: Basic Administration

<http://www.sun.com/io>

DIAGNOSTICS In addition to being logged, the following messages may also appear on the system console. All messages are formatted in the following manner:

WARNING: <device path> <usb<instance number>>: Error message...

where <instance number> is the instance number of hubd and <device path> is the physical path to the device in `/devices` directory. Messages from the root hub are displayed with a `usb<instance number>` prefix instead of `hub<instance number>` as the root hub is an integrated part of the host controller.

Connecting device on port <number> failed.

The driver failed to enumerate the device connected on port <number> of hub. If enumeration fails, disconnect and re-connect.

Use of a USB 1.0 hub behind a high speed port may cause unexpected failures.

Devices connected to a USB 1.0 hub which are in turn connected to an external USB 2.0 hub, may misbehave unexpectedly or suddenly go offline. This is due to a documented incompatibility between USB 1.0 hubs and USB 2.0 hub Transaction Translators. Please use only USB 2.0 or USB 1.1 hubs behind high-speed ports.

Connecting a high speed device to a non-high speed hub (port x) will result in a loss of performance. Please connect the device to a high speed port to get the maximum performance.

USB 2.0 devices connected to USB 1.0 or 1.1 hubs cannot run at their highest speed, even when the hub is in turn connected to a high-speed port. Reconnect without going through a USB 1.0 or 1.1 hub for best performance.

Global over current condition, please disconnect hub.

The driver detected an over current condition. This means that the aggregate current being drawn by the devices on the downstream port exceeds a preset value. Refer to section 7.2.1.2 and 11.13 of the *Universal Serial Bus Specification 2.0*. You must remove and insert this hub to render it and its downstream devices functional again. If this message continues to display for a particular hub, you may need to remove downstream devices to eliminate the problem.

Cannot access <device>. Please reconnect.

This hub has been disconnected because a device other than the original one has been inserted. The driver informs you of this fact by displaying the name of the original device.

Devices not identical to the previous one on this port. Please disconnect and reconnect.

Same condition as described above; however in this case, the driver is unable to identify the original device with a name string.

Local power has been lost, please disconnect hub.

A USB self-powered hub has lost external power. All USB devices connected down-stream from this hub will cease to function. Disconnect the hub, plug in the external power-supply and then plug in the hub again.

Hub driver supports max of <n> ports on hub. Hence, using the first <number of physical ports> of <n> ports available.

The current hub driver supports hubs that have <n> ports or less. A hub with more than <n> ports has been plugged in. Only the first <n> out of the total <number of physical ports> ports are usable.

i2o_bs(7D)

NAME	i2o_bs – Block Storage OSM for I2O
SYNOPSIS	<pre>disk@local target id#:a through u disk@local target id#:a through u raw</pre>
DESCRIPTION	<p>The I2O Block Storage OSM abstraction (BSA, which also is referred to as block storage class) layer is the primary interface that Solaris operating environments use to access block storage devices. A block storage device provides random access to a permanent storage medium. The <code>i2o_bs</code> device driver uses I2O Block Storage class messages to control the block device; and provides the same functionality (<code>ioctl</code>s, for example) that is present in the Solaris device driver like <code>'cmdk, dadk'</code> on x86 for disk. The maximum size disk supported by <code>i2o_bs</code> is the same as what is available on x86.</p> <p>The <code>i2o_bs</code> is currently implemented version 1.5 of Intelligent IO specification.</p> <p>The block files access the disk using the system's normal buffering mechanism and are read and written without regard to physical disk records. There is also a "raw" interface that provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in one I/O operation; raw I/O is therefore considerably more efficient when many bytes are transmitted. The names of the block files are found in <code>/dev/dsk</code>; the names of the raw files are found in <code>/dev/rdsk</code>.</p> <p>I2O associates each block storage device with a unique ID called a <i>local target id</i> that is assigned by I2O hardware. This information can be acquired by the block storage OSM through I2O Block Storage class messages. For Block Storage OSM, nodes are created in <code>/devices/pci#/pci#</code> which include the local target ID as one component of device name that the node refers to. However the <code>/dev</code> names and the names in <code>/dev/dsk</code> and <code>/dev/rdsk</code> do not encode the local target id in any part of the name.</p> <p>For example, you might have the following:</p> <pre>/devices/ /dev/dsk name ----- /devices/pci@0,0/pci101e,0@10,1/disk@10:a /dev/dsk/c1d0s0</pre> <p>I/O requests to the disk must have an offset and transfer length that is a multiple of 512 bytes or the driver returns an <code>EINVAL</code> error.</p> <p>Slice 0 is normally used for the root file system on a disk, slice 1 is used as a paging area (for example, swap), and slice 2 for backing up the entire <code>fdisk</code> partition for Solaris software. Other slices may be used for <code>usr</code> file systems or system reserved area.</p> <p><code>Fdisk</code> partition 0 is to access the entire disk and is generally used by the <code>fdisk(1M)</code> program.</p>
FILES	<pre>/dev/dsk/cndn [s p]n block device /dev/rdsk/cndn [s p]n raw device</pre> <p>where:</p>

i2o_bs(7D)

cn controller n
dn instance number
sn UNIX system slice n (0-15)
pn fdisk partition (0)

/kernel/drv/i2o_bs

i2o_bs driver

/kernel/drv/i2o_bs.conf

Configuration file

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO [fdisk\(1M\)](#), [format\(1M\)](#), [mount\(1M\)](#), [lseek\(2\)](#), [read\(2\)](#), [write\(2\)](#), [readdir\(3C\)](#), [vfstab\(4\)](#), [acct.h\(3HEAD\)](#), [attributes\(5\)](#), [dkio\(7I\)](#)

i2o_scsi(7D)

NAME | i2o_scsi – an I2O OS specific module that supports SCSCA interface.

DESCRIPTION | The `i2o_scsi` OSM module is a SCSI HBA driver that supports the SCSCA interface. It supports both SCSI Adapter Class and SCSI Peripheral Class functions. It translates the SCSI packet coming down from the SCSCA into an I2O SCSI Peripheral Class message, passes it along to the IOP which in turn passes it to the HDM (hardware specific module).

It also uses SCSI Adapter Class functions to manage the SCSI adapter and SCSI bus. For each SCSI Adapter Class I2O device (a SCSI controller), it claims the SCSI Peripheral class devices which are attached to that port. The existing SCSI target drivers which use the SCSCA interface should only work with `i2o_scsi`. This includes target drivers like `sd`, `st` and so on.

FILES | `/kernel/drv/i2o_scsi.conf` configuration file for the `i2o_scsi` driver; there are no user-configurable options in this file

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO | `attributes(5)`

Solaris 10 Installation Guide: Basic Installations

NAME	ib – InfiniBand Bus Nexus Driver
DESCRIPTION	The <code>ib</code> (IB nexus) driver is a pseudo nexus driver that supports enumeration of port devices, VPPA (Virtual Physical Point Attachment), HCA_SVC (HCA Service) devices, and I/O controllers (IOC) on the InfiniBand fabric that are visible to the host and provides interfaces to <code>cfgadm_ib(1M)</code> to manage hot-plugging of IB devices. The <code>ib</code> nexus driver enumerates the port device, VPPA devices and HCA_SVC devices based on entries specified in the <code>ib.conf</code> file. IOC devices are enumerated on demand. The IB nexus driver uses InfiniBand Device Manager services (<code>ibdm(7D)</code>) to enumerate port devices, VPPA devices, HCA_SVC devices, and IOCs on the IB fabric.
CONFIGURATION	You configure the <code>ib</code> driver by defining properties in the <code>ib.conf</code> file. The IB nexus driver supports the following properties:

PROPERTY NAME	DEFAULT	POSSIBLE VALUES
<code>port-svc-list</code>	""	List of service names, for example: <code>srv</code>
<code>vppa-svc-list</code>	""	List of service names, for example: <code>ipib</code>
<code>hca-svc-list</code>	""	List of service names, for example: <code>hca_nfs</code>

The `port-svc-list` property defines the list of port communication service names per port. The IB nexus driver creates a device instance for each entry in this property per Host Channel Adapter (HCA) port. The `ib.conf` file contains a `port-svc-list=""` entry by default. You update `port-svc-list` with service names you want to add to the system.

The `vppa-svc-list` property defines the list of VPPA communication service names per port per partition key. The IB nexus driver creates a device instance for each entry in this property per Host Channel Adapter (HCA) port. The `ib.conf` file contains a `vppa-svc-list=""` entry by default. You update `vppa-svc-list` with service names you want to add to the system.

The `hca-svc-list` property defines the list of HCA_SVC communication service names per HCA. The IB nexus driver creates a device instance for each entry in this property per Host Channel Adapter (HCA). The `ib.conf` file contains a `hca-svc-list=""` entry by default. You update `hca-svc-list` with service names you want to add to the system.

The service name specified in `port-svc-list`, `vppa-svc-list` and `hca-svc-list` must be unique, be a maximum of four characters long, and is limited to digits 0-9 and letters a-z and A-Z.

ib(7D)

IOC drivers (which are parented by the IB nexus driver) may themselves have .conf files. To distinguish those cases from pseudo drivers parented by IB nexus, such drivers should include the "ib-node-type" property with value "merge" in the IOC driver.conf file. That property ensures that properties from the .conf file are merged with other properties found through hardware probing.

EXAMPLES

Example 1: Default ib.conf file:

```
#
# Copyright 2001-2003 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
port-svc-list=""
vppa-svc-list="";
hca-svc-list="";
```

In Example 1, the IB nexus driver does not create any port/vppa/hca_svc device instances.

Example 2: A sample ib.conf file with one entry for "srv" service:

```
port-svc-list="srv"
vppa-svc-list="";
hca-svc-list="";
```

The IB nexus driver creates one `srv` service instance for every HCA port that exists on the host. For example, if there are two HCAs, each with two ports on the host, the IB nexus driver creates four instances of the `srv` service.

Example 3:

```
port-svc-list="srv"
vppa-svc-list="ipib";
hca-svc-list="";
```

If there are two HCAs in the system with two ports each and each port has two valid PKEY values, the IB nexus driver creates four instances of `srv` service (one for each port). It also creates eight instances of `ipd` service (one per each port/PKEY combination).

Example 4:

```
port-svc-list="srv";
vppa-svc-list="ipib";
hca-svc-list="hca_nfs";
```

The IB nexus driver creates one instance of `hca_nfs` service for each HCA in the system.

Example 5: IOC driver.conf

```
ib-node-type="merge";
enable-special-mode="on";
```

FILES /kernel/drv/ib
32-bit x86 ELF kernel module

/kernel/drv/amd64/ib
64-bit x86 ELF kernel module

/kernel/drv/sparcv9/ib
64-bit SPARC ELF kernel module

/kernel/drv/ib.conf
driver configuration file

ATTRIBUTES See `attributes(5)` for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86
Availability	SUNWhea, SUNWib
Interface Stability	Consolidation Private

SEE ALSO `cfgadm_ib(1M)`, `driver.conf(4)`, `ib(4)`, `attributes(5)`, `ibcm(7D)`, `ibdm(7D)`, `ibt1(7D)`

Writing Device Drivers

InfiniBand Architecture Specification, Volume 1: Release 1.1

System Administration Guide: Basic Administration

DIAGNOSTICS In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

ib: WARNING: Error message...

unit-address property in %s.conf not well-formed.

The `driver.conf` file does not have a valid "unit-addr" property defined. This property is an array of strings.

Cannot name node from %s.conf.

The `driver.conf` file does not have a valid "unit-addr" property defined. This property is an array of strings.

Waiting for Port %d initialization.

Waiting for port initialization from subnet manager.

ibcm(7D)

NAME ibcm – Solaris InfiniBand Communication Manager

DESCRIPTION The Solaris InfiniBand Communication Manager (IBCM) is a Solaris kernel misc module that adheres to the *InfiniBand Architecture Specification, Volume 1: Release 1.1* for InfiniBand Communication Management Class.

IBCM provides a transport layer abstraction to IB clients to set up reliable connected channels along with service, multicast, and path lookup-related functionality. IBCM implements the CM protocol as per the *InfiniBand Architecture Specification, Volume 1: Release 1.1* and utilizes the InfiniBand Management Framework module for all IB management-related functionality and the InfiniBand Transport Layer (see [ibtl\(7D\)](#)) for all IB Verbs-related functionality.

FILES /kernel/misc/ibcm
32-bit x86 ELF kernel module

/kernel/misc/amd64/ibcm
64-bit x86 ELF kernel module

/kernel/misc/sparcv9/ibcm
64-bit SPARC ELF kernel module

ATTRIBUTES See [attributes\(5\)](#) for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Consolidation Private
Availability	SUNWhea

SEE ALSO [attributes\(5\)](#), [ibtl\(7D\)](#)

InfiniBand Architecture Specification, Volume 1: Release 1.1

NAME	ibd – Infiniband IPoIB device driver
SYNOPSIS	/dev/ibd*
DESCRIPTION	<p>The ibd driver implements the IETF IP over Infiniband protocol and provides IPoIB service for all IBA ports present in the system.</p> <p>The ibd driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, dlpi(7P)). The ibd driver provides basic support for the IBA Unreliable Datagram Queue Pair hardware. Functions include QP initialization, frame transmit and receive, multicast and promiscuous mode support, and statistics reporting.</p> <p>Use the cloning, character-special device /dev/ibd to access all ibd devices installed within the system.</p> <p>The ibd driver is dependent on GLD, a loadable kernel module that provides the ibd driver with the DLPI and STREAMS functionality required of a LAN driver. Except as noted in the Application Programming Interface section of this manpage, see gld(7D) for more details on the primitives supported by the driver. The GLD module is located at /kernel/misc/sparcv9/gld on 64 bit systems and at /kernel/misc/gld on 32 bit systems.</p> <p>The ibd driver expects certain configuration of the IBA fabric prior to operation (which also implies the SM must be active and managing the fabric). Specifically, the IBA multicast group representing the IPv4 limited broadcast address 255.255.255.255 (also defined as broadcast-GID in IETF documents) must be created prior to initializing the device. IBA properties (including mtu, qkey and sl) of this group is used by the driver to create any other IBA multicast group as instructed by higher level (IP) software. The driver probes for the existence of this broadcast-GID during attach(9E).</p>
APPLICATION PROGRAMMING INTERFACE (DLPI)	<p>The values returned by the driver in the DL_INFO_ACK primitive in response to your DL_INFO_REQ are:</p> <ul style="list-style-type: none"> ■ Maximum SDU is the MTU associated with the broadcast-GID group, less the 4 byte IPoIB header. ■ Minimum SDU is 0. ■ dlsap address length is 22. ■ MAC type is DL_IB. ■ The sap length value is -2, meaning the physical address component is followed immediately by a 2-byte sap component within the DLSAP address. ■ Broadcast address value is the MAC address consisting of the 4 bytes of QPN 00:FF:FF:FF prepended to the IBA multicast address of the broadcast-GID. <p>Due to the nature of link address definition for IPoIB, the DL_SET_PHYS_ADDR_REQ DLPI primitive is not supported.</p> <p>In the transmit case for streams that have been put in raw mode via the DLIOCRAW ioctl, the DLPI application must prepend the 20 byte IPoIB destination address to the data it wants to transmit over-the-wire. In the receive case,</p>

ibd(7D)

applications receive the IP/ARP datagram along with the IETF defined 4 byte header.

WARNING

This section describes warning messages that might be generated by the driver. Please note that while the format of these messages may be modified in future versions, the same general information will be provided.

While joining IBA multicast groups corresponding to IP multicast groups as part of multicast promiscuous operations as required by IP multicast routers, or as part of running `snoop(1M)`, it is possible that joins to some multicast groups can fail due to inherent resource constraints in the IBA components. In such cases, warning message similar to the following appear in the system log, indicating the interface on which the failure occurred:

```
NOTICE: ibd0: Could not get list of IBA multicast groups
NOTICE: ibd0: IBA promiscuous mode missed multicast group
NOTICE: ibd0: IBA promiscuous mode missed new multicast gid
```

Also, if the IBA SM indicates that multicast trap support is suspended or unavailable, the system log contains a message similar to:

```
NOTICE: ibd0: IBA multicast support degraded due to
unavailability of multicast traps
```

And when the SM indicates trap support is restored:

```
NOTICE: ibd0: IBA multicast support restored due to
availability of multicast traps
```

Additionally, if the IBA link transitions to an unavailable state (that is, the IBA link state becomes "Down," "Initialize" or "Armed") and then becomes active again, the driver tries to rejoin previously joined groups if required. Failure to rejoin multicast groups will trigger messages like:

```
NOTICE: ibd0: Failure on port up to rejoin multicast gid
```

If the corresponding HCA port is in the unavailable state defined above when initializing an ibd interface using `ifconfig(1M)`, a message is emitted by the driver:

```
NOTICE: ibd0: Port is not active
```

Further, as described above, if the broadcast-GID is not found, or the associated MTU is higher than what the HCA port can support, the following messages are printed to the system log:

```
NOTICE: ibd0: IPoIB broadcast group absent
NOTICE: ibd0: IPoIB broadcast group MTU 4096 greater than port's
maximum MTU 2048
```

In all cases of these reported problems when running `ifconfig(1M)`, it should be checked that IBA cabling is intact, an SM is running on the fabric, and the broadcast-GID with appropriate properties has been created in the IBA partition.

CONFIGURATION

The IPoIB service comes preconfigured on all HCA ports in the system. To turn the service off, or back on after turning it off, refer to documentation in `cfgadm_ib(1M)`.

FILES /dev/ibd*
special character device

/kernel/drv/ib.conf
configuration file to start IPoIB service

/kernel/drv/sparcv9/ibd
64-bit SPARC device driver

/kernel/drv/amd64/ibd
64-bit x86 device driver

/kernel/drv/ibd
32-bit x86 device driver

SEE ALSO [cfgadm\(1M\)](#), [cfgadm_ib\(1M\)](#), [ifconfig\(1M\)](#), [syslogd\(1M\)](#), [gld\(7D\)](#), [ib\(7D\)](#), [kstat\(7D\)](#), [streamio\(7I\)](#), [dlpi\(7P\)](#), [attributes\(5\)](#), [attach\(9E\)](#)

ibdm(7D)

NAME | ibdm – Solaris InfiniBand Device Manager

DESCRIPTION | The Infiniband Device Manager (IBDM) is an IBTF-compliant kernel misc module. IBDM adheres to the InfiniBand Device Management class as described in *InfiniBand Architecture Specification, Volume 1: Release 1.1* and enumerates all the devices which are visible from a given host and maintains a data base of all IB devices visible to the host. IBDM provides interfaces to the IB nexus driver that enables the driver to retrieve information about IB devices on the fabric.

FILES | /kernel/misc/ibdm
32-bit x86 ELF kernel module

/kernel/misc/amd64/ibdm
64-bit x86 ELF kernel module

/kernel/misc/sparcv9/ibdm
64-bit SPARC ELF kernel module

ATTRIBUTES | See `attributes(5)` for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Contract Consolidation Private
Availability	SUNWhea

SEE ALSO | `attributes(5)`, `ib(7D)`, `ibt1(7D)`, `ibcm(7D)`
InfiniBand Architecture Specification, Volume 1: Release 1.1

DIAGNOSTICS | None.

- NAME** | ibmf – InfiniBand Management Transport Framework
- DESCRIPTION** | The InfiniBand (IB) Management Transport Framework provides the mechanisms for IB management modules to communicate with other InfiniBand management modules such as the Subnet Administration process. It also provides helper functions such as Subnet Administration Access (SAA) for commonly performed operations.
- FILES** | /kernel/misc/ibmf
32-bit ELF kernel misc module (x86 platform only).
/kernel/misc/amd64/ibmf
64-bit ELF kernel misc module (x86 platform only).
/kernel/misc/sparcv9/ibmf
64-bit ELF kernel misc module (SPARC platform only).
- ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PCI-based systems
Interface stability	Consolidation Private
Availability	SUNWib

- SEE ALSO** | `ibt1(7D)`
InfiniBand Architecture Specification, Version 1.1
www.infinibandta.org

ibtl(7D)

NAME	ibtl – Solaris InfiniBand Transport Layer						
DESCRIPTION	<p>InfiniBand (IB) is an I/O technology based on switched fabrics. The Solaris InfiniBand Transport Layer (IBTL) is a Solaris kernel misc module and adheres to the <i>IB Architecture Version 1.1</i> specification and provides a transport layer abstraction to IB client drivers.</p> <p>IBTL implements the programming interfaces for the Solaris InfiniBand Transport Framework (IBTF), consisting of the IB Channel Interface (CI) and the IB Transport Interface (TI).</p> <p>The CI consists of Host Channel Adapters (HCAs) and HCA drivers. A host is attached to the IB fabric through the CI layer. The Solaris InfiniBand CI is Sun's API rendering of the InfiniBand Architecture (IBTA) "verbs" specification.</p> <p>The Solaris InfiniBand TI is the kernel service driver interface into the Solaris InfiniBand Transport Framework. It provides transport and communications setup programming interfaces for Unreliable Datagram (UD) and Reliable Connected (RC) transport types only.</p>						
FILES	<p>/kernel/misc/ibtl 32-bit x86 ELF kernel misc module</p> <p>/kernel/misc/amd64/ibtl 64-bit x86 ELF kernel misc module</p> <p>/kernel/misc/sparcv9/ibtl 64-bit SPARC ELF kernel module</p>						
ATTRIBUTES	<p>See <code>attributes(5)</code> for a description of the following attribute:</p> <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Consolidation Private</td></tr><tr><td>Availability</td><td>SUNWhea, SUNWib</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Consolidation Private	Availability	SUNWhea, SUNWib
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Consolidation Private						
Availability	SUNWhea, SUNWib						
SEE ALSO	<p><code>attributes(5)</code>, <code>ib(7D)</code>, <code>ibcm(7D)</code>, <code>ibdm(7D)</code></p> <p><i>InfiniBand Architecture Specification, Volume 1: Release 1.1</i></p>						

NAME	icmp6 – Internet Control Message Protocol for Internet Protocol Version 6		
SYNOPSIS	<pre>#include <sys/socket.h> #include <netinet/in.h> #include <netinet/ip_icmp.h> #include <netinet/icmp6.h> s = socket(AF_INET6, SOCK_RAW, proto); t = t_open("/dev/icmp6", O_RDWR);</pre>		
DESCRIPTION	<p>The ICMP6 protocol is the error and control message protocol used with Version 6 of the Internet Protocol. It is used by the kernel to handle and report errors in protocol processing. It is also used for IPv6 neighbor and router discovery, and for multicast group membership queries and reports. It may also be accessed by programs using the socket interface or the Transport Level Interface (TLI) for network monitoring and diagnostic functions. When used with the socket interface, a “raw socket” type is used. The protocol number for ICMP6, used in the <i>proto</i> parameter to the socket call, can be obtained from <code>getprotobyname(3SOCKET)</code>. ICMP6 file descriptors and sockets are connectionless and are normally used with the <code>t_sndudata</code> / <code>t_rcvudata</code> and the <code>sendto()</code> / <code>recvfrom()</code> calls. They may also be used with the <code>sendmsg()</code> / <code>recvmsg()</code> calls when sending or receiving ancillary data.</p> <p>Outgoing packets automatically have an Internet Protocol Version 6 (IPv6) header and zero or more IPv6 extension headers prepended. These headers are prepended by the kernel. Unlike ICMP for IPv4, the <code>IP_HDRINCL</code> option is not supported for ICMP6, so ICMP6 applications neither build their own outbound IPv6 headers, nor do they receive the inbound IPv6 headers with received data. IPv6 extension headers and relevant fields of the IPv6 header may be set or received as ancillary data to a <code>sendmsg(3SOCKET)</code> or <code>recvmsg(3SOCKET)</code> system call. Each of these fields and extension headers may also be set on a per socket basis with the <code>setsockopt(3SOCKET)</code> system call. Such “sticky” options are used on all outgoing packets unless overridden by ancillary data. When any ancillary data is present with a <code>sendmsg(3SOCKET)</code> system call, all sticky options are ignored for that system call, but subsequently remain configured.</p> <p>ICMP6 is a datagram protocol layered above IPv6. Received ICMP6 messages may be reflected back to users of higher-level protocols such as TCP or UDP as error returns from system calls. A copy of each ICMP6 error message received by the system is provided to every holder of an open ICMP6 socket or TLI descriptor.</p>		
SEE ALSO	<p><code>getprotobyname(3SOCKET)</code>, <code>recv(3SOCKET)</code>, <code>recvmsg(3SOCKET)</code>, <code>send(3SOCKET)</code>, <code>sendmsg(3SOCKET)</code>, <code>setsockopt(3SOCKET)</code>, <code>t_rcvudata(3NSL)</code>, <code>t_sndudata(3NSL)</code>, inet6(7P), ip6(7P), routing(7P)</p> <p>Conta, A. and Deering, S., <i>RFC 2463, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification</i>, The Internet Society, December 1998.</p>		
DIAGNOSTICS	<p>A socket operation may fail with one of the following errors returned:</p> <table border="0"> <tr> <td style="vertical-align: top;"><code>EISCONN</code></td> <td>An attempt was made to establish a connection on a socket which already has one, or when trying to send a</td> </tr> </table>	<code>EISCONN</code>	An attempt was made to establish a connection on a socket which already has one, or when trying to send a
<code>EISCONN</code>	An attempt was made to establish a connection on a socket which already has one, or when trying to send a		

icmp6(7P)

	datagram with the destination address specified and the socket is already connected.
ENOTCONN	An attempt was made to send a datagram, but no destination address is specified, and the socket has not been connected.
ENOBUFS	The system ran out of memory for an internal data structure.
EADDRNOTAVAIL	An attempt was made to create a socket with a network address for which no network interface exists.
ENOMEM	The system was unable to allocate memory for an internal data structure.
ENOPROTOOPT	An attempt was made to set an IPv4 socket option on an IPv6 socket.
EINVAL	An attempt was made to set an invalid or malformed socket option.
EAFNOSUPPORT	An attempt was made to bind or connect to an IPv4 or mapped address, or to specify an IPv4 or mapped address as the next hop.

NAME	icmp, ICMP – Internet Control Message Protocol						
SYNOPSIS	<pre>#include <sys/socket.h> #include <netinet/in.h> #include <netinet/ip_icmp.h> s = socket(AF_INET, SOCK_RAW, proto); t = t_open("/dev/icmp", O_RDWR);</pre>						
DESCRIPTION	<p>ICMP is the error and control message protocol used by the Internet protocol family. It is used by the kernel to handle and report errors in protocol processing. It may also be accessed by programs using the socket interface or the Transport Level Interface (TLI) for network monitoring and diagnostic functions. When used with the socket interface, a “raw socket” type is used. The protocol number for ICMP, used in the <i>proto</i> parameter to the socket call, can be obtained from <code>getprotobyname(3SOCKET)</code>. ICMP file descriptors and sockets are connectionless, and are normally used with the <code>t_sndudata</code> / <code>t_rcvudata</code> and the <code>sendto()</code> / <code>recvfrom()</code> calls.</p> <p>Outgoing packets automatically have an Internet Protocol (IP) header prepended to them. Incoming packets are provided to the user with the IP header and options intact.</p> <p>ICMP is an datagram protocol layered above IP. It is used internally by the protocol code for various purposes including routing, fault isolation, and congestion control. Receipt of an ICMP “redirect” message will add a new entry in the routing table, or modify an existing one. ICMP messages are routinely sent by the protocol code. Received ICMP messages may be reflected back to users of higher-level protocols such as TCP or UDP as error returns from system calls. A copy of all ICMP message received by the system is provided to every holder of an open ICMP socket or TLI descriptor.</p>						
SEE ALSO	<p><code>getprotobyname(3SOCKET)</code>, <code>recv(3SOCKET)</code>, <code>send(3SOCKET)</code>, <code>t_rcvudata(3NSL)</code>, <code>t_sndudata(3NSL)</code>, inet(7P), ip(7P), routing(7P)</p> <p>Postel, Jon, <i>Internet Control Message Protocol — DARPA Internet Program Protocol Specification</i>, RFC 792, Network Information Center, SRI International, Menlo Park, Calif., September 1981.</p>						
DIAGNOSTICS	<p>A socket operation may fail with one of the following errors returned:</p> <table border="0"> <tr> <td style="vertical-align: top;">EISCONN</td> <td>An attempt was made to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected.</td> </tr> <tr> <td style="vertical-align: top;">ENOTCONN</td> <td>An attempt was made to send a datagram, but no destination address is specified, and the socket has not been connected.</td> </tr> <tr> <td style="vertical-align: top;">ENOBUFS</td> <td>The system ran out of memory for an internal data structure.</td> </tr> </table>	EISCONN	An attempt was made to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected.	ENOTCONN	An attempt was made to send a datagram, but no destination address is specified, and the socket has not been connected.	ENOBUFS	The system ran out of memory for an internal data structure.
EISCONN	An attempt was made to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected.						
ENOTCONN	An attempt was made to send a datagram, but no destination address is specified, and the socket has not been connected.						
ENOBUFS	The system ran out of memory for an internal data structure.						

icmp(7P)

EADDRNOTAVAIL

An attempt was made to create a socket with a network address for which no network interface exists.

NOTES

Replies to ICMP "echo" messages which are source routed are not sent back using inverted source routes, but rather go back through the normal routing mechanisms.

NAME	idn – inter-domain network device driver
SYNOPSIS	<code>/dev/idn</code>
DESCRIPTION	<p>The <code>idn</code> driver is a multi-thread, loadable, clonable, STREAMS-based pseudo driver that supports the connectionless Data Link Provider Interface <code>dlpi(7P)</code> over the Sun Enterprise 10000 Gigplane-XB Interconnect. This connection is permitted only between domains within the same Sun Enterprise 10000 server.</p> <p>The <code>idn</code> driver supports 1 to 32 logical network interfaces that can be connected to domains linked to the local domain through the <code>domain_link(1M)</code> command. (See <code>domain_link(1M)</code> in the <i>Sun Enterprise 10000 SSP 3.2 Reference Manual</i> for more information.) The <code>idn</code> driver works in conjunction with the System Service Processor (SSP) to perform domain linking/unlinking and automated linking upon host bootup.</p> <p>The <code>/dev/idn</code> device is used to access all IDN services provided by the system.</p>
IDN and DLPI	<p>The <code>idn</code> driver is a style-2 Data Link Service provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code>-type messages are interpreted as DLPI primitives. For the <code>idn</code> driver to associate the opened stream with a particular device (<code>ppa</code>), you must send an explicit <code>DL_ATTACH_REQ</code> message. The <code>ppa</code> ID is interpreted as an unsigned long and indicates the corresponding device instance (unit) number. The <code>DL_ERROR_ACK</code> error is returned by the driver if the <code>ppa</code> field value does not correspond to a valid device-instance number for the system. The device is initialized on first attach and de-initialized (stopped) on the last detach.</p> <ul style="list-style-type: none"> ■ The maximum SDU is configurable by using the <code>idn.conf</code> file and has a range of 512 bytes to 512 Kbytes. The default value is 16384 bytes. ■ The minimum SDU is 0. ■ The Service Access Pointer (SAP) address length is 8. ■ The MAC type is <code>DL_ETHER</code>. ■ The SAP length value is -2, meaning the physical address component is followed immediately by a 2-byte SAP component within the DLSAP address. ■ The service mode is <code>DL_CLDLS</code>. ■ Optional quality of service (QOS) is not presently supported; accordingly, the QOS fields are 0. ■ The provider style is <code>DL_STYLE2</code>. ■ The version is <code>DL_VERSION_2</code>. ■ The broadcast address value is Ethernet/IEEE broadcast address (0xFFFFFFFF). The <code>idn</code> driver supports broadcast by issuing messages to each target individually. The <code>idn</code> driver is inherently a point-to-point network between domains. When the <code>idn</code> driver is in the <code>DL_ATTACHED</code> state, the user must send a <code>DL_BIND_REQ</code> request to associate a particular SAP with the stream. The <code>idn</code> driver interprets the SAP field within the <code>DL_BIND_REQ</code> message as an Ethernet type and valid values for the SAP field are in the range of 0 to 0xFFFF. Only one Ethernet type can be bound to the stream at any time.

idn(7d)

If a SAP with a value of 0 is selected, the receiver will be in 802.3 mode. All frames received from the media having a type field in the range of 0 to 1500 are assumed to be 802.3 frames and are routed up all open streams which are bound to SAP value 0. If more than one stream is in 802.3 mode, then the frame will be duplicated and routed up as multiple stream DL_UNITDATA_IND messages.

In transmission, the driver checks the SAP field of the DL_BIND_REQ to determine if the SAP value is 0, and if the destination type field is in the range of 0 to 1500. If either is true, the driver computes the length of the message, (excluding the initial message block M_PROTO mb1k) of all subsequent DL_UNITDATA_REQ messages and transmits 802.3 frames that have this value in the MAC frame header length field.

The driver also supports raw M_DATA mode. When the user sends a DLIOCRAW ioctl, the particular stream is put in raw mode. A complete frame and a proper ether header is expected as part of the data.

The DLSAP address format consists of the 6-byte, physical address component (Ethernet) followed immediately by the 2-byte SAP component (type), producing an 8-byte DLSAP address. Applications should *not* hardcode to this particular implementation-specific DLSAP address format, but instead should use information returned in the DL_INFO_ACK primitive to compose and decompose DLSAP addresses. The SAP length, full DLSAP length, and SAP physical ordering are included within the DL_INFO_ACK primitive. The physical address length can be computed by subtracting the SAP length from the full DLSAP address length or by issuing the DL_PHYS_ADDR_REQ message to obtain the current physical address associated with the stream.

When the idn driver is in the DL_BOUND state, you can transmit frames on the IDN by sending DL_UNITDATA_REQ messages to the driver. The driver then routes received IDN frames up the open and bound streams having a SAP which matches the Ethernet type as DL_UNITDATA_IND messages. If necessary, received IDN frames are duplicated and routed up multiple open streams. The DLSAP address contained within the DL_UNITDATA_REQ and DL_UNITDATA_IND messages consists of both the SAP (type) and physical (Ethernet) components.

IDN Primitives

In addition to the mandatory connectionless DLPI message set, the idn driver supports the following primitives:

The DL_ENABMULTI_REQ and DL_DISABMULTI_REQ primitives which enable or disable, respectively, the reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. These primitives are accepted by the driver in any state following the DL_ATTACHED state.

The DL_PROMISCON_REQ and DL_PROMISCOFF_REQ primitives, which with the DL_PROMISC_PHYS flag set in the dl_level field, enable or disable, respectively, the reception of all promiscuous frames on the media, including frames generated by the local domain. When used with the DL_PROMISC_SAP flag set in the dl_level field,

idn(7d)

these primitives enable or disable, respectively, the reception of all SAP (Ethernet type) values. When used with the `DL_PROMISC_MULTI` flag set in the `d1_level` field, these primitives enable or disable, respectively, the reception of all multicast group addresses. The effect of each is always on a per-stream basis and independent of the other SAP and physical level configurations on this stream or other streams.

The `DL_PHYS_ADDR_REQ` primitive which returns the 6-octet, Ethernet address associated with (or attached to) the stream in the `DL_PHYS_ADDR_ACK` primitive. This primitive is valid only in states following a successful `DL_ATTACH_REQ` request.

Because the driver maintains domain address information in the address to direct packets to the correct destination, the `DL_SET_PHYS_ADDR_REQ` primitive is not allowed.

FILES The following files are supported:

`/dev/idn`

IDN special character device

`/platform/SUNW,Ultra-Enterprise-10000/kernel/drv/idn.conf`

System-wide and per-interface default device driver properties

SEE ALSO `netstat(1M)`, `ndd(1M)`, `dlpi(7P)`

`domain_link(1M)` in the *Sun Enterprise 10000 SSP 3.2 Reference Manual*.

Sun Enterprise 10000 InterDomain Networks User Guide

NOTES The `idn` driver supports a set of properties that can be set by using the `driver.conf` file for the IDN. See the *Sun Enterprise 10000 InterDomain Networks User Guide* for more information about the properties in the `driver.conf(4)`, (`idn.conf`, for IDNs).

ifb(7d)

NAME	ifb – 24-bit PCI color frame buffer and graphics accelerator driver
DESCRIPTION	<p>The <code>ifb</code> driver is the device driver for the Sun Expert3D graphics accelerators. The Expert3D is a high resolution, high performance PCI graphics framebuffer providing hardware texture mapping. The Expert3D also supports 1920x1200 double buffered, z-buffered display and 1280 x 1024 stereo.</p> <p>The <code>ifbdaemon</code> process loads the <code>ifb</code> microcode at system startup time and during the resume sequence of a suspend-resume cycle.</p>
FILES	<p><code>/dev/fbs/ifbn</code> Device special file</p> <p><code>/usr/lib/ifb.unicode</code> ifb microcode</p> <p><code>/usr/sbin/ifbdaemon</code> ifb microcode loader</p>
SEE ALSO	<code>SUNwifb_config(1M)</code>

NAME	ifp – ISP2100 Family Fibre Channel Host Bus Adapter Driver						
SYNOPSIS	PCI SUNW,ifp@pci-slot						
DESCRIPTION	<p>The ifp Host Bus Adapter is a SCSI compliant nexus driver for the Qlogic ISP2100/ISP2100A chips. These chips support Fibre Channel Protocol for SCSI on Private Fibre Channel Arbitrated loops.</p> <p>The ifp driver interfaces with SCSI disk target driver, ssd(7D), and the SCSI-3 Enclosure Services driver, ses(7D). Only SCSI devices of type <code>disk</code> and <code>ses</code> are supported at present time.</p> <p>The ifp driver supports the standard functions provided by the SCSI interface. It supports auto request sense (cannot be turned off) and tagged queueing by default. The driver requires that all devices have unique hard addresses defined by switch settings in hardware. Devices with conflicting hard addresses will not be accessible.</p>						
FILES	<table border="0"> <tr> <td>/kernel/drv/ifp</td> <td>ELF Kernel Module</td> </tr> <tr> <td>/kernel/drv/sparcv9/ifp</td> <td>ELF Kernel Module (64-bit version)</td> </tr> <tr> <td>/kernel/drv/ifp.conf</td> <td>Driver configuration file</td> </tr> </table>	/kernel/drv/ifp	ELF Kernel Module	/kernel/drv/sparcv9/ifp	ELF Kernel Module (64-bit version)	/kernel/drv/ifp.conf	Driver configuration file
/kernel/drv/ifp	ELF Kernel Module						
/kernel/drv/sparcv9/ifp	ELF Kernel Module (64-bit version)						
/kernel/drv/ifp.conf	Driver configuration file						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SPARC</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SPARC		
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SPARC						
SEE ALSO	<p><code>luxadm(1M)</code>, <code>prtconf(1M)</code>, <code>driver.conf(4)</code>, <code>attributes(5)</code>, <code>ses(7D)</code>, <code>ssd(7D)</code></p> <p><i>Writing Device Drivers</i>,</p> <p><i>ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL)</i>,</p> <p><i>ANSI X3.269-1996, Fibre Channel Protocol for SCSI (FCP)</i>,</p> <p><i>ANSI X3.270-1996, SCSI-3 Architecture Model (SAM)</i>,</p> <p><i>Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)</i>,</p> <p><i>ISP2100 Firmware Interface Specification, QLogic Corporation</i></p>						
DIAGNOSTICS	<p>The messages described below are some that may appear on the system console, as well as being logged.</p> <p>This first set of messages may be displayed while the ifp driver is initially trying to attach. All of these messages mean that the ifp driver was unable to attach. These messages are preceded by "ifp<number>", where "<number>" is the instance number of the ISP2100 Host Bus Adapter.</p>						

ifp(7D)

Device is using a hilevel intr, unused

The device was configured with an interrupt level that cannot be used with this ifp driver. Check the device.

Failed to alloc soft state

Driver was unable to allocate space for the internal state structure. Driver did not attach to device; SCSI devices will be inaccessible.

Bad soft state

Driver requested an invalid internal state structure. Driver did not attach to device; SCSI devices will be inaccessible.

Unable to map pci config registers

Unable to map biu registers

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot alloc tran

Driver was unable to obtain a transport handle to be able to communicate with SCSA framework. Driver did not attach to device; SCSI devices will be inaccessible.

ddi_create_minor_node failed

Driver was unable to create devctl minor node that is used by luxadm(1M) for administering the loop. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot alloc dma handle

Driver was unable allocate a dma handle for communicating with the Host Bus Adapter. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot alloc cmd area

Driver was unable allocate dma memory for request and response queues. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot bind cmd area

Driver was unable to bind dma handle to the cmd area. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot alloc fcal handle

Driver was unable allocate a dma handle for retrieving loop map from the Host Bus Adapter. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot bind portdb

Driver was unable to bind fcal port handle to the memory used for obtaining port database. Driver did not attach to device; SCSI devices will be inaccessible.

scsi_hba_attach failed

Driver was unable to attach to the SCSA framework. Driver did not attach to device; SCSI devices will be inaccessible.

Unable to create hotplug thread

Driver was not able to create the kernel thread used for hotplug support. Driver did not attach to device; SCSI devices will be inaccessible.

Cannot add intr

Driver was not able to add the interrupt routine to the kernel. Driver did not attach to device; SCSI devices will be inaccessible.

Unable to attach

Driver was unable to attach to the hardware for some reason that may be printed. Driver did not attach to device; SCSI devices will be inaccessible.

The following set of messages may be display at any time. They will be printed with the full device pathname followed by the shorter form described above.

Firmware checksum incorrect

Firmware has an invalid checksum and will not be downloaded.

Chip reset timeout

ISP chip failed to reset in the time allocated; may be bad hardware.

Stop firmware failed

Stopping the firmware failed; may be bad hardware.

Load ram failed

Unable to download new firmware into the ISP chip.

DMA setup failed

The DMA setup failed in the host adapter driver on a `scsi_pkt`. This will return `TRAN_BADPKT` to a SCSA target driver.

Bad request pkt type**Bad request pkt****Bad request pkt hdr****Bad req pkt order**

The ISP Firmware rejected the packet as being set up incorrectly. This will cause the ifp driver to call the target completion routine with the reason of `CMD_TRAN_ERR` set in the `scsi_pkt`. Check the target driver for correctly setting up the packet.

Firmware error

The ISP chip encountered a firmware error of some kind. This error will cause the ifp driver to do error recovery by resetting the chip.

DMA Failure (event)

The ISP chip encountered a DMA error while reading from the request queue (event is 8003) or writing to the response queue (event is 8004). This error will cause the ifp driver to do error recovery by resetting the chip.

Fatal error, resetting interface

This is an indication that the ifp driver is doing error recovery. This will cause all outstanding commands that have been transported to the ifp driver to be completed via the `scsi_pkt` completion routine in the target driver with reason of `CMD_RESET` and status of `STAT_BUS_RESET` set in the `scsi_pkt`.

target t, duplicate port wwns

The driver detected target *t* to be having the same port WWN as a different target; this is not supposed to happen. Target *t* will become inaccessible.

ifp(7D)

target *t*, duplicate switch settings

The driver detected devices with the same switch setting *t*. All such devices will become inaccessible.

WWN changed on target *t*

The World Wide Name (WWN) has changed on the device with switch setting *t*.

target *t*, unknown device type *dt*

The driver does not know the device type *dt* reported by the device with switch setting *t*.

NAME	if_tcp, if – general properties of Internet Protocol network interfaces
DESCRIPTION	A network interface is a device for sending and receiving packets on a network. It is usually a hardware device, although it can be implemented in software. Network interfaces used by the Internet Protocol (IPv4 or IPv6) must be STREAMS devices conforming to the Data Link Provider Interface (DLPI). See dlpi(7P) .
APPLICATION PROGRAMMING INTERFACE	An interface becomes available to IP when it is opened and the IP module is pushed onto the stream with the <code>I_PUSH</code> <code>ioctl(2)</code> command. (See streamio(7I)). The <code>SIOCSLIFNAME</code> <code>ioctl(2)</code> is issued to specify the name of the interface and to indicate whether it is IPv4 or IPv6. This may be initiated by the kernel at boot time or by a user program after the system is running. Each interface must be assigned an IP address with the <code>SIOCSLIFADDR</code> <code>ioctl()</code> before it can be used. On interfaces where the network-to-link layer address mapping is static, only the network number is taken from the <code>ioctl()</code> request; the remainder is found in a hardware specific manner. On interfaces which provide dynamic network-to-link layer address mapping facilities (for example, 10Mb/s Ethernets using arp(7P)), the entire address specified in the <code>ioctl()</code> is used. A routing table entry for destinations on the network of the interface is installed automatically when an interface's address is set.
IOCTLS	The following <code>ioctl()</code> calls may be used to manipulate IP network interfaces. Unless specified otherwise, the request takes an <code>lifreq</code> structure as its parameter. This structure has the form: <pre> /* Interface request structure used for socket ioctls. All */ /* interface ioctls must have parameter definitions which */ /* begin with lifr_name. The remainder may be interface specific. */ struct lifreq { #define LIFNAMSIZ 32 char lifr_name[LIFNAMSIZ]; /* if name, for example "le1" */ union { int lifru_addrlen; /* for subnet/token etc */ uint_t lifru_ppa; /* SIOCSLIFNAME */ } lifr_lifru1; union { struct sockaddr_storage lifru_addr; struct sockaddr_storage lifru_dstaddr; struct sockaddr_storage lifru_broadaddr; struct sockaddr_storage lifru_token; /* With lifr_addrlen */ struct sockaddr_storage lifru_subnet; /* With lifr_addrlen */ int lifru_index; /* interface index */ uint64_t lifru_flags; /* SIOC?LIFFLAGS */ int lifru_metric; uint_t lifru_mtu; char lifru_data[1]; /* intfce.dep. data/ char lifru_enaddr[6]; int lifr_muxid[2]; /* mux id's for arp and ip */ struct lif_nd_req lifru_nd_req; struct lif_ifinfo_req lifru_ifinfo_req; zoneid_t lifru_zone; /* SIOC[GS]LIFZONE */ } lifr_lifru; #define lifr_addrlen lifr_lifru1.lifru_addrlen #define lifr_ppa lifr_lifru1.lifru_ppa /* Driver's ppa */ </pre>

if_tcp(7P)

```

#define lifr_addr      lifr_lifru.lifru_addr      /* address */
#define lifr_dstaddr  lifr_lifru.lifru_dstaddr
#define lifr_broadaddr lifr_lifru.lifru_broadaddr /* broadcast address */
#define lifr_token    lifr_lifru.lifru_token    /* address token */
#define lifr_subnet   lifr_lifru.lifru_subnet   /* subnet prefix */
#define lifr_index    lifr_lifru.lifru_index    /* interface index */
#define lifr_flags    lifr_lifru.lifru_flags    /* flags */
#define lifr_metric   lifr_lifru.lifru_metric   /* metric */
#define lifr_mtu      lifr_lifru.lifru_mtu      /* mtu */
#define lifr_data     lifr_lifru.lifru_data
#define lifr_enaddr   lifr_lifru.lifru_enaddr   /* ethernet address */
#define lifr_index    lifr_lifru.lifru_index    /* interface index */
#define lifr_ip_muxid lifr_lifru.lif_muxid[0]
#define lifr_arp_muxid lifr_lifru.lif_muxid[1]
#define lifr_nd       lifr_lifru.lifru_nd_req   /* SIOCLIF*ND */
#define lifr_ifinfo   lifr_lifru.lifru_ifinfo_req /* SIOC[GS]LIFLNKINFO */
#define lifr_zone     lifr_lifru.lifru_zone     /* SIOC[GS]LIFZONE */
};

```

SIOCSLIFADDR	Set interface address. Following the address assignment, the “initialization” routine for the interface is called.
SIOCGLIFADDR	Get interface address.
SIOCSLIFDSTADDR	Set point to point address for interface.
SIOCGLIFDSTADDR	Get point to point address for interface.
SIOCSLIFFLAGS	Set interface flags field. If the interface is marked down, any processes currently routing packets through the interface are notified.
SIOCGLIFFLAGS	Get interface flags.
SIOCGLIFCONF	Get interface configuration list. This request takes an <code>lifconf</code> structure (see below) as a value-result parameter. The <code>lifc_len</code> field should be initially set to the size of the buffer pointed to by <code>lifc_buf</code> . On return it will contain the length, in bytes, of the configuration list. The <code>lifc_family</code> field should be set to <code>AF_UNSPEC</code> to retrieve both <code>AF_INET</code> and <code>AF_INET6</code> interfaces. The <code>lifc_flags</code> field should be initially set to zero.
SIOCGLIFNUM	Get number of interfaces. This request returns an integer which is the number of interface descriptions (<code>struct lifreq</code>) that will be returned by the <code>SIOCGLIFCONF</code> ioctl; that is, it gives an indication of how large <code>lifc_len</code> has to be. This request takes an <code>lifnum</code> structure (see below) as a value-result parameter. The <code>lifn_family</code> field should be set to <code>AF_UNSPEC</code> to count both <code>AF_INET</code> and <code>AF_INET6</code> interfaces. The <code>lifn_flags</code> field should be initially set to zero.

SIOCSLIFMTU	Set the maximum transmission unit (MTU) size for interface. Place the result of this request in <code>lifru_mtu</code> field. The MTU can not exceed the physical MTU limitation (which is reported in the DLPI <code>DL_INFO_ACK</code> message).
SIOCGLIFMTU	Get the maximum transmission unit size for interface. Place the result of this request in <code>ifru_mtu</code> field.
SIOCSLIFMETRIC	Set the metric associated with the interface. The metric is used by routing daemons such as <code>in.routed(1M)</code> .
SIOCGLIFMETRIC	Get the metric associated with the interface.
SIOCGLIFMUXID	Get the <code>ip</code> and <code>arp muxid</code> associated with the interface.
SIOCSLIFMUXID	Set the <code>ip</code> and <code>arp muxid</code> associated with the interface.
SIOCGLIFINDEX	Get the interface index associated with the interface.
SIOCSLIFINDEX	Set the interface index associated with the interface.
SIOCGLIFZONE	Get the zone associated with the interface.
SIOCSLIFZONE	Set the zone associated with the interface.
SIOCLIFADDIF	Add a new logical interface on a physical interface using an unused logical unit number. If the physical interface is part of an IP multipathing group, the logical interface may be added to a different physical interface in the same group. Upon return, the <code>lifr_name</code> field contains the name of the actual logical interface created.
SIOCLIFREMOVEIF	Remove a logical interface by specifying its IP address or logical interface name. When the IP address is specified and the interface is part of an IP multipathing group, the logical interface is removed from the physical interface in the group which holds the IP address.
SIOCSLIFTOKEN	Set the address token used to form IPv6 link-local addresses and for stateless address autoconfiguration.
SIOCGLIFTOKEN	Get the address token used to form IPv6 link-local addresses and for stateless address autoconfiguration.
SIOCSLIFSUBNET	Set the subnet prefix associated with the interface.
SIOCGLIFSUBNET	Get the subnet prefix associated with the interface.
SIOCSLIFLNKINFO	Set link specific parameters for the interface.

if_tcp(7P)

<code>SIOCGLIFLNKINFO</code>	Get link specific parameters for the interface.
<code>SIOCLIFDELND</code>	Delete a neighbor cache entry for IPv6.
<code>SIOCLIFGETND</code>	Get a neighbor cache entry for IPv6.
<code>SIOCLIFSETND</code>	Set a neighbor cache entry for IPv6.
<code>SIOCTMYADDR</code>	Test if the address is assigned to this node. This request takes an <code>sioc_addrreq</code> structure (see below) as a value-result parameter. The <code>sa_addr</code> field should be set to the address to test. The <code>sa_res</code> field will contain a non-zero value if the address is assigned to this node.
<code>SIOCSSLIFUSESRC</code>	Set the interface from which to choose a source address. The <code>lifr_index</code> field has the interface index corresponding to the interface whose address is to be used as the source address for packets going out on the interface whose name is provided by <code>lifr_name</code> . If the <code>lifr_index</code> field is set to zero, the previous setting is cleared. See <code>ifconfig(1M)</code> for examples of the 'usesrc' option.
<code>SIOCGLIFUSESRC</code>	Get the interface index of the interface whose address is used as the source address for packets going out on the interface provided by <code>lifr_name</code> field. The value is retrieved in the <code>lifr_index</code> field. See <code>ifconfig(1M)</code> for examples of the 'usesrc' option.
<code>SIOCGLIFSRCOF</code>	Get the interface configuration list for interfaces that use an address hosted on the interface provided by the <code>lifs_ifindex</code> field in the <code>lifsrcof</code> struct (see below), as a source address. The application sets <code>lifs_maxlen</code> to the size (in bytes) of the buffer it has allocated for the data. On return, the kernel sets <code>lifs_len</code> to the actual size required. Note, the application could set <code>lifs_maxlen</code> to zero to query the kernel of the required buffer size instead of estimating a buffer size. The application tests ' <code>lifs_len <= lifs_maxlen</code> ' -- if that's true, the buffer was big enough and the application has an accurate list. If it is false, it needs to allocate a bigger buffer and try again, and <code>lifs_len</code> provides a hint of how big to make the next trial. See <code>ifconfig(1M)</code> for examples of the 'usesrc' option. The <code>lifsrcof</code> structure has the form: /* * Structure used in SIOCGLIFSRCOF to get the interface * configuration list for those interfaces that use an address * hosted on the interface (set in lifs_ifindex), as the source

```

        * address.
        */
struct lifsrcof {
    uint_t  lifs_ifindex; /* addr on this interface used as the
    size_t  lifs_maxlen;  /* size of buffer: input */
    size_t  lifs_len;     /* size of buffer: output */
    union {
        caddr_t lifsu_buf;
        struct  lifreq *lifsu_req;
    } lifs_lifsu;
#define lifs_buf lifs_lifsu.lifsu_buf /* buffer address */
#define lifs_req lifs_lifsu.lifsu_req /* array returned */
};

```

SIOCTONLINK

Test if the address is directly reachable, for example, that it can be reached without going through a router. This request takes an `sioc_addrreq` structure (see below) as a value-result parameter. The `sa_addr` field should be set to the address to test. The `sa_res` field will contain a non-zero value if the address is onlink.

SIOCTMYSITE

Test if the address is part of the same site as this node. This request takes an `sioc_addrreq` structure (see below) as a value-result parameter. The `sa_addr` field should be set to the address to test. The `sa_res` field will contain a non-zero value if the address is in the same site.

The `lifconf` structure has the form:

```

/*
 * Structure used in SIOCGLIFCONF request.
 * Used to retrieve interface configuration
 * for machine (useful for programs which
 * must know all networks accessible).
 */
struct lifconf {
    sa_family_t  lifc_family;
    int          lifc_flags;          /* request specific interfaces */
    int          lifc_len;           /* size of associated buffer */
    union {
        caddr_t  lifcu_buf;
        struct  ifreq *lifcu_req;
    } lifc_lifcu;
#define lifc_buf lifc_lifcu.lifcu_buf /* buffer address */
#define lifc_req lifc_lifcu.lifcu_req /* array of structures returned */
};

```

The `sioc_addrreq` structure has the form:

```

/* Structure used in SIOCGLIFNUM request. */
struct lifnum {
    sa_family_t  lifn_family;

```

if_tcp(7P)

```
    int          lifn_flags; /* request specific interfaces */
    int          lifn_count; /* Result */
};

/*
 * Argument structure for SIOCT* address testing ioctls.
 */
struct sioc_addrreq {
    struct sockaddr_storage sa_addr; /* Address to test */
    int                    sa_res; /* Result - 0/1 */
};
```

The following `ioctl()` calls are maintained for compatibility but only apply to IPv4 network interfaces, since the data structures are too small to hold an IPv6 address. Unless specified otherwise, the request takes an `ifreq` structure as its parameter. This structure has the form:

```
/* Interface request structure used for socket ioctls. All */
/* interface ioctls must have parameter definitions which */
/* begin with lifr_name. The remainder may be interface specific. */
struct ifreq {
#define IFNAMSIZ 16
    char    lifr_name[IFNAMSIZ]; /* if name, for example */
                                     /* "le1" */

    union {
        struct sockaddr    ifru_addr;
        struct sockaddr    ifru_dstaddr;
        char    ifru_ename[IFNAMSIZ]; /* other if name */
        struct sockaddr    ifru_broadaddr;
        short    ifru_flags;
        int    ifru_metric;
        char    ifru_data[1]; /* intfce. depen. data */
        char    ifru_enaddr[6];
        int    if_muxid[2]; /* mux id's for arp and ip */
        int    ifru_index; /* interface index */
    } ifr_ifru;

#define ifr_addr    ifr_ifru.ifru_addr /* address */
#define ifr_dstaddr    ifr_ifru.ifru_dstaddr /* other end of p-to-p link */
#define ifr_ename    ifr_ifru.ifru_ename /* other if name */
#define ifr_broadaddr    ifr_ifru.ifru_broadaddr /* broadcast address */
#define ifr_flags    ifr_ifru.ifru_flags /* flags */
#define ifr_index    ifr_ifru.ifru_index /* interface index */
#define ifr_metric    ifr_ifru.ifru_metric /* metric */
#define ifr_data    ifr_ifru.ifru_data /* for use by interface */
#define ifr_enaddr    ifr_ifru.ifru_enaddr /* ethernet address */
};
```

SIOCSIFADDR Set interface address. Following the address assignment, the "initialization" routine for the interface is called.

SIOCGIFADDR Get interface address.

SIOCSIFDSTADDR	Set point to point address for interface.
SIOCGIFDSTADDR	Get point to point address for interface.
SIOCSIFFLAGS	Set interface flags field. If the interface is marked down, any processes currently routing packets through the interface are notified.
SIOCGIFFLAGS	Get interface flags.
SIOCGIFCONF	Get interface configuration list. This request takes an <code>ifconf</code> structure (see below) as a value-result parameter. The <code>ifc_len</code> field should be initially set to the size of the buffer pointed to by <code>ifc_buf</code> . On return it will contain the length, in bytes, of the configuration list.
SIOCGIFNUM	Get number of interfaces. This request returns an integer which is the number of interface descriptions (<code>struct ifreq</code>) that will be returned by the <code>SIOCGIFCONF</code> ioctl; that is, it gives an indication of how large <code>ifc_len</code> has to be.
SIOCSIFMTU	Set the maximum transmission unit (MTU) size for interface. Place the result of this request in <code>ifru_metric</code> field. The MTU has to be smaller than physical MTU limitation (which is reported in the DLPI <code>DL_INFO_ACK</code> message).
SIOCGIFMTU	Get the maximum transmission unit size for interface. Place the result of this request in <code>ifru_metric</code> field.
SIOCSIFMETRIC	Set the metric associated with the interface. The metric is used by routine daemons such as <code>in.routed(1M)</code> .
SIOCGIFMETRIC	Get the metric associated with the interface.
SIOCGIFMUXID	Get the <code>ip</code> and <code>arp muxid</code> associated with the interface.
SIOCSIFMUXID	Set the <code>ip</code> and <code>arp muxid</code> associated with the interface.
SIOCGIFINDEX	Get the interface index associated with the interface.
SIOCSIFINDEX	Set the interface index associated with the interface.

The `ifconf` structure has the form:

```
/*
 * Structure used in SIOCGIFCONF request.
 * Used to retrieve interface configuration
 * for machine (useful for programs which
 * must know all networks accessible).
 */
```

if_tcp(7P)

```
struct ifconf {
    int    ifc_len;                /* size of associated buffer */
    union {
        caddr_t    ifcu_buf;
        struct ifreq    *ifcu_req;
    } ifc_ifcu;

#define ifc_buf    ifc_ifcu.ifcu_buf    /* buffer address */
#define ifc_req    ifc_ifcu.ifcu_req    /* array of structures returned */
};
```

IFF_Flags

To display addressing information for each interface, enter `ifconfig -a`. IFF_flags are defined in `/usr/include/net/if.h`. The flags are:

```
#define IFF_UP            0x00000001    /* Interface is up */
#define IFF_BROADCAST    0x00000002    /* Broadcast address valid */
#define IFF_DEBUG        0x00000004    /* Turn on debugging */
#define IFF_LOOPBACK     0x00000008    /* Loopback net */
#define IFF_POINTOPOINT  0x00000010    /* Interface is point-to-point link */
#define IFF_NOTRAILERS   0x00000020    /* Avoid use of trailers */
#define IFF_RUNNING      0x00000040    /* Resources allocated */
#define IFF_NOARP        0x00000080    /* No address resolution protocol */
#define IFF_PROMISC      0x00000100    /* Receive all packets */
#define IFF_ALLMULTI     0x00000200    /* Receive all multicast packets */
#define IFF_INTELLIGENT  0x00000400    /* Protocol code on board */
#define IFF_MULTICAST    0x00000800    /* Supports multicast */
#define IFF_MULTI_BCAST  0x00001000    /* Multicast using broadcast address */
#define IFF_UNNUMBERED   0x00002000    /* Non-unique address */
#define IFF_DHCPRUNNING  0x00004000    /* DHCP controls this interface */
#define IFF_PRIVATE      0x00008000    /* Do not advertise */

#define IFF_NOXMIT        0x00010000    /* Do not transmit packets */
#define IFF_NOLOCAL      0x00020000    /* No address - just on-link subnet */
#define IFF_DEPRECATED   0x00040000    /* Interface address deprecated */
#define IFF_ADDRCONF     0x00080000    /* Address from stateless addrconf */

#define IFF_ROUTER        0x00100000    /* Router on this interface */
#define IFF_NONUD         0x00200000    /* No NUD on this interface */
#define IFF_ANYCAST      0x00400000    /* Anycast address */
#define IFF_NORTEXCH     0x00800000    /* Do not exchange routing info */

#define IFF_IPV4          0x01000000    /* IPv4 interface */
#define IFF_IPV6          0x02000000    /* IPv6 interface */
#define IFF_MIPRUNNING    0x04000000    /* Mobile IP controls this interface */

#define IFF_NOFAILOVER    0x008000000    /* Don't failover on NIC failure */

#define IFF_FAILED        0x010000000    /* NIC has failed */
#define IFF_STANDBY       0x020000000    /* Standby NIC-use on failures */
#define IFF_INACTIVE      0x040000000    /* Standby active or not ? */
#define IFF_OFFLINE       0x080000000    /* NIC has been offlined */

#define IFF_XRESOLV       0x100000000    /* IPv6 external resolver */
#define IFF_COS_ENABLED   0x200000000    /* If CoS marking is supported */
#define IFF_VIRTUAL       0x200000000    /* Does not send or receive packets */
```

For more information on IFF_ flags, see [ifconfig\(1M\)](#).

ERRORS	EPERM	The effective user id of the calling process is not superuser.
	ENXIO	The <code>lifr_name</code> member of the <code>lifreq</code> structure contains an invalid value. For <code>SIOCGLIFSRCOF</code> , the <code>lifs_ifindex</code> member of the <code>lifsrcof</code> structure contains an invalid value. For <code>SIOCSLIFUSESRC</code> , this error is returned if the <code>lifr_index</code> is set to an invalid value.
	EBADADDR	Wrong address family or malformed address.
	EBUSY	For <code>SIOCSLIFFLAGS</code> , this error is returned when the order of bringing the primary/physical interface (for example, <code>eri0</code>) and a secondary/logical interface associated with the same physical interface (for example, <code>eri0:1</code>) up or down is violated. The physical interface must be configured up first and cannot be configured down until all the corresponding logical interfaces have been configured down.
	EINVAL	For <code>SIOCSLIFMTU</code> , this error is returned when the requested MTU size is invalid. This error indicates the MTU size is greater than the MTU size supported by the DLPI provider or less than 68 (for IPv4) or less than 1200 (for IPv6).] For <code>SIOCSLIFUSESRC</code> , this error is returned if either the <code>lifr_index</code> or <code>lifr_name</code> identify interfaces that are already part of an existing IPMP group.
	EEXIST	For <code>SIOCLIFADDIF</code> , this error is returned if the <code>lifr_name</code> member in the <code>lifreq</code> structure corresponds to an interface that already has the PPA specified by <code>lifru_ppa</code> plumbed.

SEE ALSO [ifconfig\(1M\)](#), [in.routed\(1M\)](#), [ioctl\(2\)](#), [streamio\(7I\)](#), [arp\(7P\)](#), [dlpi\(7P\)](#), [ip\(7P\)](#), [ip6\(7P\)](#)

inet6(7P)

NAME	inet6 – Internet protocol family for Internet Protocol version 6
SYNOPSIS	<pre>#include <sys/types.h> #include <netinet/in.h></pre>
DESCRIPTION	<p>The <code>inet6</code> protocol family implements a collection of protocols that are centered around the Internet Protocol version 6 (IPv6) and share a common address format. The <code>inet6</code> protocol family can be accessed using the socket interface, where it supports the <code>SOCK_STREAM</code>, <code>SOCK_DGRAM</code>, and <code>SOCK_RAW</code> socket types, or the Transport Level Interface (TLI), where it supports the connectionless (<code>T_CLTS</code>) and connection oriented (<code>T_COTS_ORD</code>) service types.</p>
PROTOCOLS	<p>The Internet protocol family for IPv6 included the Internet Protocol Version 6 (IPv6), the Neighbor Discovery Protocol (NDP), the Internet Control Message Protocol (ICMPv6), the Transmission Control Protocol (TCP), and the User Datagram Protocol (UDP).</p> <p>TCP supports the socket interface's <code>SOCK_STREAM</code> abstraction and TLI's <code>T_COTS_ORD</code> service type. UDP supports the <code>SOCK_DGRAM</code> socket abstraction and the TLI <code>T_CLTS</code> service type. See tcp(7P) and udp(7P). A direct interface to IPv6 is available using the socket interface. See ip6(7P). ICMPv6 is used by the kernel to handle and report errors in protocol processing. It is also accessible to user programs. See icmp6(7P). NDP is used to translate 128-bit IPv6 addresses into 48-bit Ethernet addresses.</p> <p>IPv6 addresses come in three types: unicast, anycast, and multicast. A unicast address is an identifier for a single network interface. An anycast address is an identifier for a set of interfaces; a packet sent to an anycast address is delivered to the "nearest" interface identified by that address, pursuant to the routing protocol's measure of distance. A multicast address is an identifier for a set of interfaces; a packet sent to a multicast address is delivered to all interfaces identified by that address. There are no broadcast addresses as such in IPv6; their functionality is superseded by multicast addresses.</p> <p>For IPv6 addresses, there are three scopes within which unicast addresses are guaranteed to be unique. The scope is indicated by the address prefix. The three varieties are link-local (the address is unique on that physical link), site-local (the address is unique within that site), and global (the address is globally unique).</p> <p>The three highest order bits for global unicast addresses are set to 001. The ten highest order bits for site-local addresses are set to 1111 1110 11. The ten highest order bits for link-local addresses are set to 1111 1110 11. For multicast addresses, the eight highest order bits are set to 1111 1111. Anycast addresses have the same format as unicast addresses.</p> <p>IPv6 addresses do not follow the concept of "address class" seen in IP.</p> <p>A global unicast address is divided into the following segments:</p> <ul style="list-style-type: none">■ The first three bits are the Format Prefix identifying a unicast address.

- The next 13 bits are the Top-Level Aggregation (TLA) identifier. For example, the identifier could specify the ISP.
- The next eight bits are reserved for future use.
- The next 24 bits are the Next-Level Aggregation (NLA) identifier.
- The next 16 bits are the Site-Level Aggregation (SLA) identifier.
- The last 64 bits are the interface ID. This will most often be the hardware address of the link in IEEE EUI-64 format.

Link-local unicast addresses are divided in this manner:

- The first ten bits are the Format Prefix identifying a link-local address.
- The next 54 bits are zero.
- The last 64 bits are the interface ID. This will most often be the hardware address of the link in IEEE EUI-64 format.

Site-local unicast addresses are divided in this manner:

- The first ten bits are the Format Prefix identifying a site-local address.
- The next 38 bits are zero.
- The next 16 bits are the subnet ID.
- The last 64 bits are the interface ID. This will most often be the hardware address of the link in IEEE EUI-64 format.

ADDRESSING

IPv6 addresses are sixteen byte quantities, stored in network byte order. The socket API uses the `sockaddr_in6` structure when passing IPv6 addresses between an application and the kernel. The `sockaddr_in6` structure has the following members:

```
sa_family_t    sin6_family;
in_port_t      sin6_port;
uint32_t       sin6_flowinfo;
struct in6_addr sin6_addr;
uint32_t       sin6_scope_id;
uint32_t       __sin6_src_id; Library routines are provided to manipulate
structures of this form. See inet(3SOCKET).
```

The `sin6_addr` field of the `sockaddr_in6` structure specifies a local or remote IPv6 address. Each network interface has one or more IPv6 addresses configured, that is, a link-local address, a site-local address, and one or more global unicast IPv6 addresses. The special value of all zeros may be used on this field to test for "wildcard" matching. Given in a `bind(3SOCKET)` call, this value leaves the local IPv6 address of the socket unspecified, so that the socket will receive connections or messages directed at any of the valid IPv6 addresses of the system. This can prove useful when a process neither knows nor cares what the local IPv6 address is, or when a process wishes to receive requests using all of its network interfaces.

inet6(7P)

The `sockaddr_in6` structure given in the `bind()` call must specify an `in6_addr` value of either all zeros or one of the system's valid IPv6 addresses. Requests to bind any other address will elicit the error `EADDRNOTAVAIL`. When a `connect(3SOCKET)` call is made for a socket that has a wildcard local address, the system sets the `sin6_addr` field of the socket to the IPv6 address of the network interface through which the packets for that connection are routed.

The `sin6_port` field of the `sockaddr_in6` structure specifies a port number used by TCP or UDP. The local port address specified in a `bind()` call is restricted to be greater than `IPPORT_RESERVED` (defined in `<netinet/in.h>`) unless the creating process is running as the super-user, providing a space of protected port numbers. In addition, the local port address cannot be in use by any socket of the same address family and type. Requests to bind sockets to port numbers being used by other sockets return the error `EADDRINUSE`. If the local port address is specified as 0, the system picks a unique port address greater than `IPPORT_RESERVED`. A unique local port address is also selected when a socket which is not bound is used in a `connect(3SOCKET)` or `sendto()` call. See `send(3SOCKET)`. This allows programs that do not care which local port number is used to set up TCP connections by simply calling `socket(3SOCKET)` and then `connect(3SOCKET)`, and then sending UDP datagrams with a `socket()` call followed by a `sendto()` call.

Although this implementation restricts sockets to unique local port numbers, TCP allows multiple simultaneous connections involving the same local port number so long as the remote IPv6 addresses or port numbers are different for each connection. Programs may explicitly override the socket restriction by setting the `SO_REUSEADDR` socket option with `setsockopt()`. See `getsockopt(3SOCKET)`.

In addition, the same port may be bound by two separate sockets if one is an IP socket and the other an IPv6 socket.

TLI applies somewhat different semantics to the binding of local port numbers. These semantics apply when Internet family protocols are used using the TLI.

SOURCE ADDRESS SELECTION

IPv6 source address selection is done on a per destination basis, and utilizes a list of rules from which the best source address is selected from candidate addresses. The candidate set comprises a set of local addresses assigned on the system which are up and not anycast. If just one candidate exists in the candidate set, it is selected.

Conceptually, each selection rule prefers one address over another, or determines their equivalence. If a rule produces a tie, a subsequent rule is used to break the tie.

The sense of some rules may be reversed on a per-socket basis using the `IPV6_SRC_PREFERENCES` socket option (see [ip6\(7P\)](#)). The flag values for this option are defined in `<netinet/in.h>` and are referenced in the description of the appropriate rules below.

As the selection rules indicate, the candidate addresses are SA and SB and the destination is D.

Prefer the same address

If $SA == D$, prefer SA. If $SB == D$, prefer SB.

Prefer appropriate scope

Here, $Scope(X)$ is the scope of X according to the IPv6 Addressing Architecture.

If $Scope(SA) < Scope(SB)$: If $Scope(SA) < Scope(D)$, then prefer SB and otherwise prefer SA.

If $Scope(SB) < Scope(SA)$: If $Scope(SB) < Scope(D)$, then prefer SA and otherwise prefer SB.

Avoid deprecated addresses

If one of the addresses is deprecated (IFF_DEPRECATED) and the other is not, prefer the one that isn't deprecated.

Prefer preferred addresses

If one of the addresses is preferred (IFF_PREFERRED) and the other is not, prefer the one that is preferred.

Prefer outgoing interface

If one of the addresses is assigned to the interface that will be used to send packets to D and the other is not, then prefer the former.

Prefer matching label

This rule uses labels which are obtained through the IPv6 default address selection policy table. See `ipaddrsel(1M)` for a description of the default contents of the table and how the table is configured.

If $Label(SA) == Label(D)$ and $Label(SB) != Label(D)$, then prefer SA.

If $Label(SB) == Label(D)$ and $Label(SA) != Label(D)$, then prefer SB.

Prefer public addresses

This rule prefers public addresses over temporary addresses, as defined in *RFC 3041*. Temporary addresses are disabled by default and may be enabled by appropriate settings in `ndpd.conf(4)`.

The sense of this rule may be set on a per-socket basis using the `IPV6_SRC_PREFERENCES` socket option. Passing the flag `IPV6_PREFER_SRC_TMP` or `IPV6_PREFER_SRC_PUBLIC` will cause temporary or public addresses to be preferred, respectively, for that particular socket. See `ip6(7P)` for more information about IPv6 socket options.

Use longest matching prefix.

This rule prefers the source address that has the longer matching prefix with the destination. Because this is the last rule and because both source addresses could have equal matching prefixes, this rule does an `xor` of each source address with the destination, then selects the source address with the smaller `xor` value in order to break any potential tie.

If $SA \wedge D < SB \wedge D$, then prefer SA.

If $SB \wedge D < SA \wedge D$, then prefer SB.

inet6(7P)

Applications can override this algorithm by calling `bind(3SOCKET)` and specifying an address.

SEE ALSO `ioctl(2)`, `bind(3SOCKET)`, `connect(3SOCKET)`, `getipnodebyaddr(3SOCKET)`, `getipnodebyname(3SOCKET)`, `getprotobyname(3SOCKET)`, `getservbyname(3SOCKET)`, `getsockopt(3SOCKET)`, `inet(3SOCKET)`, `send(3SOCKET)`, [icmp6\(7P\)](#), [ip6\(7P\)](#), [tcp\(7P\)](#), [udp\(7P\)](#)

Conta, A. and Deering, S., *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*, RFC 1885, December 1995.

Deering, S. and Hinden, B., *Internet Protocol, Version 6 (IPv6) Specification*, RFC 1883, December 1995.

Hinden, B. and Deering, S., *IP Version 6 Addressing Architecture*, RFC 1884, December 1995.

Draves, R., *RFC 3484, Default Address Selection for IPv6*. The Internet Society. February 2003.

Narten, T., and Draves, R. *RFC 3041, Privacy Extensions for Stateless Address Autoconfiguration in IPv6*. The Internet Society. January 2001.

NOTES The IPv6 support is subject to change as the Internet protocols develop. Users should not depend on details of the current implementation, but rather the services exported.

NAME	inet – Internet protocol family				
SYNOPSIS	<pre>#include <sys/types.h> #include <netinet/in.h></pre>				
DESCRIPTION	<p>The Internet protocol family implements a collection of protocols which are centered around the Internet Protocol (“IP”) and which share a common address format. The Internet family protocols can be accessed using the socket interface, where they support the <code>SOCK_STREAM</code>, <code>SOCK_DGRAM</code>, and <code>SOCK_RAW</code> socket types, or the Transport Level Interface (TLI), where they support the connectionless (<code>T_CLTS</code>) and connection oriented (<code>T_COTS_ORD</code>) service types.</p>				
PROTOCOLS	<p>The Internet protocol family is comprised of the Internet Protocol (“IP”), the Address Resolution Protocol (“ARP”), the Internet Control Message Protocol (“ICMP”), the Transmission Control Protocol (“TCP”), and the User Datagram Protocol (“UDP”).</p> <p>TCP supports the socket interface’s <code>SOCK_STREAM</code> abstraction and TLI’s <code>T_COTS_ORD</code> service type. UDP supports the <code>SOCK_DGRAM</code> socket abstraction and the TLI <code>T_CLTS</code> service type. See tcp(7P) and udp(7P). A direct interface to IP is available using both TLI and the socket interface (see ip(7P)). ICMP is used by the kernel to handle and report errors in protocol processing. It is also accessible to user programs (see icmp(7P)). ARP is used to translate 32-bit IP addresses into 48-bit Ethernet addresses. See arp(7P).</p> <p>The 32-bit IP address is divided into network number and host number parts. It is frequency-encoded. The most-significant bit is zero in Class A addresses, in which the high-order 8 bits represent the network number. Class B addresses have their high order two bits set to 10 and use the high-order 16 bits as the network number field. Class C addresses have a 24-bit network number part of which the high order three bits are 110. Sites with a cluster of IP networks may chose to use a single network number for the cluster; this is done by using subnet addressing. The host number portion of the address is further subdivided into subnet number and host number parts. Within a subnet, each subnet appears to be an individual network. Externally, the entire cluster appears to be a single, uniform network requiring only a single routing entry. Subnet addressing is enabled and examined by the following <code>ioctl(2)</code> commands. They have the same form as the <code>SIOCSIFADDR</code> command.</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top;"><code>SIOCSIFNETMASK</code></td> <td>Set interface network mask. The network mask defines the network part of the address; if it contains more of the address than the address type would indicate, then subnets are in use.</td> </tr> <tr> <td style="vertical-align: top;"><code>SIOCGIFNETMASK</code></td> <td>Get interface network mask.</td> </tr> </table>	<code>SIOCSIFNETMASK</code>	Set interface network mask. The network mask defines the network part of the address; if it contains more of the address than the address type would indicate, then subnets are in use.	<code>SIOCGIFNETMASK</code>	Get interface network mask.
<code>SIOCSIFNETMASK</code>	Set interface network mask. The network mask defines the network part of the address; if it contains more of the address than the address type would indicate, then subnets are in use.				
<code>SIOCGIFNETMASK</code>	Get interface network mask.				
ADDRESSING	<p>IP addresses are four byte quantities, stored in network byte order. IP addresses should be manipulated using the byte order conversion routines. See <code>byteorder(3SOCKET)</code>.</p> <p>Addresses in the Internet protocol family use the <code>sockaddr_in</code> structure, which has that following members:</p>				

inet(7P)

```
short    sin_family;
ushort_t sin_port;
struct   in_addr sin_addr;
char     sin_zero[8];
```

Library routines are provided to manipulate structures of this form; See `inet(3SOCKET)`.

The `sin_addr` field of the `sockaddr_in` structure specifies a local or remote IP address. Each network interface has its own unique IP address. The special value `INADDR_ANY` may be used in this field to effect “wildcard” matching. Given in a `bind(3SOCKET)` call, this value leaves the local IP address of the socket unspecified, so that the socket will receive connections or messages directed at any of the valid IP addresses of the system. This can prove useful when a process neither knows nor cares what the local IP address is or when a process wishes to receive requests using all of its network interfaces. The `sockaddr_in` structure given in the `bind(3SOCKET)` call must specify an `in_addr` value of either `INADDR_ANY` or one of the system’s valid IP addresses. Requests to bind any other address will elicit the error `EADDRNOTAVAIL`. When a `connect(3SOCKET)` call is made for a socket that has a wildcard local address, the system sets the `sin_addr` field of the socket to the IP address of the network interface that the packets for that connection are routed through.

The `sin_port` field of the `sockaddr_in` structure specifies a port number used by TCP or UDP. The local port address specified in a `bind(3SOCKET)` call is restricted to be greater than `IPPORT_RESERVED` (defined in `<<netinet/in.h>>`) unless the creating process is running as the superuser, providing a space of protected port numbers. In addition, the local port address must not be in use by any socket of same address family and type. Requests to bind sockets to port numbers being used by other sockets return the error `EADDRINUSE`. If the local port address is specified as 0, then the system picks a unique port address greater than `IPPORT_RESERVED`. A unique local port address is also picked when a socket which is not bound is used in a `connect(3SOCKET)` or `sendto` (see `send(3SOCKET)`) call. This allows programs which do not care which local port number is used to set up TCP connections by simply calling `socket(3SOCKET)` and then `connect(3SOCKET)`, and to send UDP datagrams with a `socket(3SOCKET)` call followed by a `sendto()` call.

Although this implementation restricts sockets to unique local port numbers, TCP allows multiple simultaneous connections involving the same local port number so long as the remote IP addresses or port numbers are different for each connection. Programs may explicitly override the socket restriction by setting the `SO_REUSEADDR` socket option with `setsockopt` (see `getsockopt(3SOCKET)`).

TLI applies somewhat different semantics to the binding of local port numbers. These semantics apply when Internet family protocols are used using the TLI.

SEE ALSO

`ioctl(2)`, `bind(3SOCKET)`, `byteorder(3SOCKET)`, `connect(3SOCKET)`, `gethostbyname(3NSL)`, `getnetbyname(3SOCKET)`, `getprotobyname(3SOCKET)`, `getservbyname(3SOCKET)`, `getsockopt(3SOCKET)`, `send(3SOCKET)`, `socket(3SOCKET)`, `arp(7P)`, `icmp(7P)`, `ip(7P)`, `tcp(7P)`, `udp(7P)`

Network Information Center, *DDN Protocol Handbook* (3 vols.), Network Information Center, SRI International, Menlo Park, Calif., 1985.

NOTES The Internet protocol support is subject to change as the Internet protocols develop. Users should not depend on details of the current implementation, but rather the services exported.

ip6(7P)

NAME	ip6 – Internet Protocol Version 6		
SYNOPSIS	<pre>#include <sys/socket.h> #include <netinet/in.h>, #include <netinet/ip6.h> s = socket(AF_INET6, SOCK_RAW, proto); t = t_open ("/dev/rawip6", O_RDWR);</pre>		
DESCRIPTION	<p>The IPv6 protocol is the next generation of the internetwork datagram delivery protocol of the Internet protocol family. Programs may use IPv6 through higher-level protocols such as the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP), or may interface directly to IPv6. See tcp(7P) and udp(7P). Direct access may be by means of the socket interface, using a “raw socket,” or by means of the Transport Level Interface (TLI). The protocol options and IPv6 extension headers defined in the IPv6 specification may be set in outgoing datagrams.</p>		
APPLICATION PROGRAMMING INTERFACE	<p>The STREAMS driver <code>/dev/rawip6</code> is the TLI transport provider that provides raw access to IPv6.</p> <p>Raw IPv6 sockets are connectionless and are normally used with the <code>sendto()</code> and <code>recvfrom()</code> calls (see <code>send(3SOCKET)</code> and <code>recv(3SOCKET)</code>), although the <code>connect(3SOCKET)</code> call may also be used to fix the destination for future datagrams. In this case, the <code>read(2)</code> or <code>recv(3SOCKET)</code> and <code>write(2)</code> or <code>send(3SOCKET)</code> calls may be used. Ancillary data may also be sent or received over raw IPv6 sockets using the <code>sendmsg(3SOCKET)</code> and <code>recvmsg(3SOCKET)</code> system calls.</p> <p>Unlike raw IP, IPv6 applications do not include a complete IPv6 header when sending; there is no IPv6 analog to the <code>IP_IP_HDRINCL</code> socket option. IPv6 header values may be specified or received as ancillary data to a <code>sendmsg(3SOCKET)</code> or <code>recvmsg(3SOCKET)</code> system call, or may be specified as “sticky” options on a per-socket basis by using the <code>setsockopt(3SOCKET)</code> system call. Such sticky options are applied to all outbound packets unless overridden by ancillary data. If any ancillary data is specified in a <code>sendmsg(3SOCKET)</code> call, all sticky options not explicitly overridden revert to default values for that datagram only; the sticky options persist as set for subsequent datagrams.</p> <p>Since <code>sendmsg(3SOCKET)</code> is not supported for <code>SOCK_STREAM</code> upper level protocols such as TCP, ancillary data is unsupported for TCP. Sticky options, however, are supported.</p> <p>Since <code>sendmsg(3SOCKET)</code> is supported for <code>SOCK_DGRAM</code> upper level protocols, both ancillary data and sticky options are supported for UDP, ICMP6, and raw IPv6 sockets.</p> <p>The socket options supported at the IPv6 level are:</p> <table><tr><td><code>IPV6_BOUND_IF</code></td><td>Limit reception transmission of packets to this interface. Takes an integer as an argument; the integer is the selected interface index.</td></tr></table>	<code>IPV6_BOUND_IF</code>	Limit reception transmission of packets to this interface. Takes an integer as an argument; the integer is the selected interface index.
<code>IPV6_BOUND_IF</code>	Limit reception transmission of packets to this interface. Takes an integer as an argument; the integer is the selected interface index.		

IPV6_UNSPEC_SRC	Boolean. Allow/disallow sending with a zero source address.
IPV6_UNICAST_HOPS	Default hop limit for unicast datagrams. This option takes an integer as an argument. Its value becomes the new default value for <code>ip6_hops</code> that IPv6 will use on outgoing unicast datagrams sent from that socket. The initial default is 60.
IPV6_CHECKSUM	Specify the integer offset in bytes into the user data of the checksum location. Does not apply to the ICMP6 protocol. Note: checksums are required for all IPv6 datagrams; this is different from IP, in which datagram checksums were optional. IPv6 will compute the ULP checksum if the value in the checksum field is zero.
IPV6_SEC_OPT	Enable or obtain IPsec security settings for this socket. For more details on the protection services of IPsec, see ipsec(7P) .
IPV6_DONTFRAG	Boolean. Control fragmentation.
IPV6_USE_MIN_MTU	Boolean. Use the minimum MTU for transmitting traffic.
IPV6_V6ONLY	Boolean. If set, only V6 packets can be sent or received
IPV6_SRC_PREFERENCES	Enable or obtain Source Address Selection rule settings for this socket. For more details on the Source Address Selection rules, see inet6(7P) .

The following options are boolean switches controlling the reception of ancillary data:

IPV6_RECVPKTINFO	Enable/disable receipt of the index of the interface the packet arrived on, and of the inbound packet's destination address.
IPV6_RECVHOPLIMIT	Enable/disable receipt of the inbound packet's current hoplimit.
IPV6_RECVHOPOPTS	Enable/disable receipt of the inbound packet's IPv6 hop-by-hop extension header.
IPV6_RECVDSTOPTS	Enable/disable receipt of the inbound packet's IPv6 destination options extension header.
IPV6_RECVRTHDR	Enable/disable receipt of the inbound packet's IPv6 routing header.
IPV6_RECVRTHDRDSTOPTS	Enable/disable receipt of the inbound packet's intermediate-hops options extension header. This option is obsolete. <code>IPV6_RECVDSTOPTS</code> turns on receipt of both destination option headers.

ip6(7P)

IPV6_RECVTCLASS	Enable/disable receipt of the traffic class of the inbound packet.
IPV6_RECVPATHMTU	Enable/disable receipt of the path mtu of the inbound packet.
The following options may be set as sticky options with <code>setsockopt(3SOCKET)</code> or as ancillary data to a <code>sendmsg(3SOCKET)</code> system call:	
IPV6_PKTINFO	Set the source address and/or interface out which the packet(s) will be sent. Takes a <code>struct ip6_pktinfo</code> as the parameter.
IPV6_HOPLIMIT	Set the initial hoplimit for outbound datagrams. Takes an integer as the parameter. Note: This option sets the hoplimit only for ancillary data or sticky options and does not change the default hoplimit for the socket; see <code>IPV6_UNICAST_HOPS</code> and <code>IPV6_MULTICAST_HOPS</code> to change the socket's default hoplimit.
IPV6_NEXTHOP	Specify the IPv6 address of the first hop, which must be a neighbor of the sending host. Takes a <code>struct sockaddr_in6</code> as the parameter. When this option specifies the same address as the destination IPv6 address of the datagram, this is equivalent to the existing <code>SO_DONTROUTE</code> option.
IPV6_HOPOPTS	Specify one or more hop-by-hop options. Variable length. Takes a complete IPv6 hop-by-hop options extension header as the parameter.
IPV6_DSTOPTS	Specify one or more destination options. Variable length. Takes a complete IPv6 destination options extension header as the parameter.
IPV6_RTHDR	Specify the IPv6 routing header. Variable length. Takes a complete IPv6 routing header as the parameter. Currently, only type 0 routing headers are supported.
IPV6_RTHDRDSTOPTS	Specify one or more destination options for all intermediate hops. May be configured, but will not be applied unless an IPv6 routing header is also configured. Variable length. Takes a complete IPv6 destination options extension header as the parameter.
IPV6_PATHMTU	Get the path MTU associated with a connected socket. Takes a <code>ip6_mtuinfo</code> as the parameter.
IPV6_TCLASS	Set the traffic class associated with outgoing packets. The parameter is an integer. If the parameter is less than -1 or greater than 256, <code>EINVAL</code> is returned. If the parameter is equal to -1, use the default. If the

parameter is between 0 and 255 inclusive, use that value.

The following options affect the socket's multicast behavior:

IPV6_JOIN_GROUP	Join a multicast group. Takes a <code>struct ipv6_mreq</code> as the parameter; the structure contains a multicast address and an interface index.
IPV6_LEAVE_GROUP	Leave a multicast group. Takes a <code>struct ipv6_mreq</code> as the parameter; the structure contains a multicast address and an interface index.
IPV6_MULTICAST_IF	The outgoing interface for multicast packets. This option takes an integer as an argument; the integer is the interface index of the selected interface.
IPV6_MULTICAST_HOPS	Default hop limit for multicast datagrams. This option takes an integer as an argument. Its value becomes the new default value for <code>ip6_hops</code> that IPv6 will use on outgoing multicast datagrams sent from that socket. The initial default is 1.
IPV6_MULTICAST_LOOP	Loopback for multicast datagrams. Normally multicast datagrams are delivered to members on the sending host. Setting the unsigned character argument to 0 will cause the opposite behavior.

The multicast socket options can be used with any datagram socket type in the IPv6 family.

At the socket level, the socket option `SO_DONTROUTE` may be applied. This option forces datagrams being sent to bypass routing and forwarding by forcing the IPv6 `hoplimit` field to 1, meaning that the packet will not be forwarded by routers.

Raw IPv6 datagrams can also be sent and received using the TLI connectionless primitives.

Datagrams flow through the IPv6 layer in two directions: from the network *up* to user processes and from user processes *down* to the network. Using this orientation, IPv6 is layered *above* the network interface drivers and *below* the transport protocols such as UDP and TCP. The Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) is logically a part of IPv6. See [icmp6\(7P\)](#).

Unlike IP, IPv6 provides no checksum of the IPv6 header. Also unlike IP, upper level protocol checksums are required. IPv6 will compute the ULP/data portion checksum if the checksum field contains a zero (see `IPV6_CHECKSUM` option above).

ip6(7P)

IPv6 extension headers in received datagrams are processed in the IPv6 layer according to the protocol specification. Currently recognized IPv6 extension headers include hop-by-hop options header, destination options header, routing header (currently, only type 0 routing headers are supported), and fragment header.

By default, the IPv6 layer will not forward IPv6 packets that are not addressed to it. This behavior can be overridden by using `routedm(1M)` to enable the `ipv6-forwarding` option. IPv6 forwarding is configured at boot time based on the setting of `routedm(1M)`'s `ipv6-forwarding` option.

For backwards compatibility, IPv6 forwarding can be enabled or disabled using `ndd(1M)`'s `ip_forwarding` variable. It is set to 1 if IPv6 forwarding is enabled, or 0 if it is disabled.

Additionally, finer-grained forwarding can be configured in IPv6. Each interface can be configured to forward IPv6 packets by setting the `IFF_ROUTER` interface flag. This flag can be set and cleared using `ifconfig(1M)`'s `router` and `-router` options. If an interface's `IFF_ROUTER` flag is set, packets can be forwarded to or from the interface. If it is clear, packets will neither be forwarded from this interface to others, nor forwarded to this interface. Setting the `ip6_forwarding` variable sets all of the IPv6 interfaces' `IFF_ROUTER` flags.

For backwards compatibility, each interface creates an `<ifname>ip6_forwarding/dev/ip6` variable that can be modified using `ndd(1M)`. An interface's `:ip6_forwarding` `ndd` variable is a boolean variable that mirrors the status of its `IFF_ROUTER` interface flag. It is set to 1 if the flag is set, or 0 if it is clear. This interface specific `<ifname>:ip6_forwarding` `ndd` variable is obsolete and may be removed in a future release of Solaris. The `ifconfig(1M)` `router` and `-router` interfaces are preferred.

The IPv6 layer will send an ICMP6 message back to the source host in many cases when it receives a datagram that can not be handled. A "time exceeded" ICMP6 message will be sent if the `ip6_hops` field in the IPv6 header drops to zero in the process of forwarding a datagram. A "destination unreachable" message will be sent by a router or by the originating host if a datagram can not be sent on because there is no route to the final destination; it will be sent by a router when it encounters a firewall prohibition; it will be sent by a destination node when the transport protocol (that is, TCP) has no listener. A "packet too big" message will be sent by a router if the packet is larger than the MTU of the outgoing link (this is used for Path MTU Discovery). A "parameter problem" message will be sent if there is a problem with a field in the IPv6 header or any of the IPv6 extension headers such that the packet cannot be fully processed.

The IPv6 layer supports fragmentation and reassembly. Datagrams are fragmented on output if the datagram is larger than the maximum transmission unit (MTU) of the network interface. Fragments of received datagrams are dropped from the reassembly queues if the complete datagram is not reconstructed within a short time period.

Errors in sending discovered at the network interface driver layer are passed by IPv6 back up to the user process.

SEE ALSO [svcs\(1\)](#), [ndd\(1M\)](#), [routeadm\(1M\)](#), [svcadm\(1M\)](#), [read\(2\)](#), [write\(2\)](#), [bind\(3SOCKET\)](#), [connect\(3SOCKET\)](#), [getsockopt\(3SOCKET\)](#), [recv\(3SOCKET\)](#), [recvmsg\(3SOCKET\)](#), [send\(3SOCKET\)](#), [sendmsg\(3SOCKET\)](#), [setsockopt\(3SOCKET\)](#), [defaultrouter\(4\)](#), [smf\(5\)](#), [icmp6\(7P\)](#), [if_tcp\(7P\)](#), [ipsec\(7P\)](#), [inet6\(7P\)](#), [routing\(7P\)](#), [tcp\(7P\)](#), [udp\(7P\)](#)

Deering, S. and Hinden, B. *RFC 2460, Internet Protocol, Version 6 (IPv6) Specification*. The Internet Society. December, 1998.

Stevens, W., and Thomas, M. *RFC 2292, Advanced Sockets API for IPv6*. Network Working Group. February 1998.

DIAGNOSTICS A socket operation may fail with one of the following errors returned:

EACCES	A <code>bind()</code> operation was attempted with a “reserved” port number and the effective user ID of the process was not the privileged user.
EADDRINUSE	A <code>bind()</code> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
EADDRNOTAVAIL	A <code>bind()</code> operation was attempted for an address that is not configured on this machine.
EINVAL	A <code>sendmsg()</code> operation with a non-NULL <code>msg_accrights</code> was attempted.
EINVAL	A <code>getsockopt()</code> or <code>setsockopt()</code> operation with an unknown socket option name was given.
EINVAL	A <code>getsockopt()</code> or <code>setsockopt()</code> operation was attempted with the IPv6 option field improperly formed; an option field was shorter than the minimum value or longer than the option buffer provided; the value in the option field was invalid.
EISCONN	A <code>connect()</code> operation was attempted on a socket on which a <code>connect()</code> operation had already been performed, and the socket could not be successfully disconnected before making the new connection.
EISCONN	A <code>sendto()</code> or <code>sendmsg()</code> operation specifying an address to which the message should be sent was attempted on a socket on which a <code>connect()</code> operation had already been performed.

ip6(7P)

EMSGSIZE	A <code>send()</code> , <code>sendto()</code> , or <code>sendmsg()</code> operation was attempted to send a datagram that was too large for an interface, but was not allowed to be fragmented (such as broadcasts).
ENETUNREACH	An attempt was made to establish a connection via <code>connect()</code> , or to send a datagram by means of <code>sendto()</code> or <code>sendmsg()</code> , where there was no matching entry in the routing table; or if an ICMP "destination unreachable" message was received.
ENOTCONN	A <code>send()</code> or <code>write()</code> operation, or a <code>sendto()</code> or <code>sendmsg()</code> operation not specifying an address to which the message should be sent, was attempted on a socket on which a <code>connect()</code> operation had not already been performed.
ENOBUFS	The system ran out of memory for fragmentation buffers or other internal data structures.
ENOMEM	The system was unable to allocate memory for an IPv6 socket option or other internal data structures.
ENOPROTOOPT	An IP socket option was attempted on an IPv6 socket, or an IPv6 socket option was attempted on an IP socket.

NOTES Applications using the sockets API must use the Advanced Sockets API for IPv6 (RFC 2292) to see elements of the inbound packet's IPv6 header or extension headers.

The `ip6` service is managed by the service management facility, `smf(5)`, under the service identifier:

```
svc:/network/initial:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using `svcadm(1M)`. The service's status can be queried using the `svcs(1)` command.

NAME	ip, IP – Internet Protocol		
SYNOPSIS	<pre>#include <sys/socket.h> #include <netinet/in.h> s = socket(AF_INET, SOCK_RAW, proto); t = t_open ("/dev/rawip", O_RDWR);</pre>		
DESCRIPTION	<p>IP is the internetwork datagram delivery protocol that is central to the Internet protocol family. Programs may use IP through higher-level protocols such as the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP), or may interface directly to IP. See tcp(7P) and udp(7P). Direct access may be by means of the socket interface, using a “raw socket,” or by means of the Transport Level Interface (TLI). The protocol options defined in the IP specification may be set in outgoing datagrams.</p> <p>Packets sent to or from this system may be subject to IPsec policy; see ipsec(7P) for more information.</p>		
APPLICATION PROGRAMMING INTERFACE	<p>The STREAMS driver <code>/dev/rawip</code> is the TLI transport provider that provides raw access to IP.</p> <p>Raw IP sockets are connectionless and are normally used with the <code>sendto()</code> and <code>recvfrom()</code> calls (see <code>send(3SOCKET)</code> and <code>recv(3SOCKET)</code>), although the <code>connect(3SOCKET)</code> call may also be used to fix the destination for future datagram. In this case, the <code>read(2)</code> or <code>recv(3SOCKET)</code> and <code>write(2)</code> or <code>send(3SOCKET)</code> calls may be used. If <code>proto</code> is <code>IPPROTO_RAW</code> or <code>IPPROTO_IGMP</code>, the application is expected to include a complete IP header when sending. Otherwise, that protocol number will be set in outgoing datagrams and used to filter incoming datagrams and an IP header will be generated and prepended to each outgoing datagram. In either case, received datagrams are returned with the IP header and options intact.</p> <p>If an application uses <code>IP_HDRINCL</code> and provides the IP header contents, the IP stack does not modify the following supplied fields under any conditions: Type of Service, DF Flag, Protocol, and Destination Address. The IP Options and IHL fields are set by use of <code>IP_OPTIONS</code>, and Total Length is updated to include any options. Version is set to the default. Identification is chosen by the normal IP ID selection logic. The source address is updated if none was specified and the TTL is changed if the packet has a broadcast destination address. Since an application cannot send down fragments (as IP assigns the IP ID), Fragment Offset is always 0. The IP Checksum field is computed by IP. None of the data beyond the IP header are changed, including the application-provided transport header.</p> <p>The socket options supported at the IP level are:</p> <table border="0"> <tr> <td style="vertical-align: top;"><code>IP_OPTIONS</code></td> <td>IP options for outgoing datagrams. This socket option may be used to set IP options to be included in each outgoing datagram. IP options to be sent are set with <code>setsockopt()</code> (see <code>getsockopt(3SOCKET)</code>). The</td> </tr> </table>	<code>IP_OPTIONS</code>	IP options for outgoing datagrams. This socket option may be used to set IP options to be included in each outgoing datagram. IP options to be sent are set with <code>setsockopt()</code> (see <code>getsockopt(3SOCKET)</code>). The
<code>IP_OPTIONS</code>	IP options for outgoing datagrams. This socket option may be used to set IP options to be included in each outgoing datagram. IP options to be sent are set with <code>setsockopt()</code> (see <code>getsockopt(3SOCKET)</code>). The		

ip(7P)

	<p><code>getsockopt(3SOCKET)</code> call returns the IP options set in the last <code>setsockopt()</code> call. IP options on received datagrams are visible to user programs only using raw IP sockets. The format of IP options given in <code>setsockopt()</code> matches those defined in the IP specification with one exception: the list of addresses for the source routing options must include the first-hop gateway at the beginning of the list of gateways. The first-hop gateway address will be extracted from the option list and the size adjusted accordingly before use. IP options may be used with any socket type in the Internet family.</p>
IP_SEC_OPT	Enable or obtain IPsec security settings for this socket. For more details on the protection services of IPsec, see ipsec(7P) .
IP_ADD_MEMBERSHIP	Join a multicast group.
IP_DROP_MEMBERSHIP	Leave a multicast group.
<p>These options take a <code>struct ip_mreq</code> as the parameter. The structure contains a multicast address which has to be set to the CLASS-D IP multicast address, and an interface address. Normally the interface address is set to <code>INADDR_ANY</code> which causes the kernel to choose the interface to join on.</p>	
IP_MULTICAST_IF	The outgoing interface for multicast packets. This option takes a <code>struct in_addr</code> as an argument, and it selects that interface for outgoing IP multicast packets. If the address specified is <code>INADDR_ANY</code> , it will use the unicast routing table to select the outgoing interface (which is the default behavior).
IP_MULTICAST_TTL	Time to live for multicast datagrams. This option takes an unsigned character as an argument. Its value is the TTL that IP will use on outgoing multicast datagrams. The default is 1.
IP_MULTICAST_LOOP	Loopback for multicast datagrams. Normally multicast datagrams are delivered to members on the sending host (or sending zone). Setting the unsigned character argument to 0 causes the opposite behavior, meaning that when multiple zones are present, the datagrams are delivered to all zones except the sending zone.
IP_RECVIF	Receive the inbound interface index.
IP_TOS	This option takes an integer argument as its input value. The least significant 8 bits of the value are used to set the Type Of Service field in the IP header of the outgoing packets.

The multicast socket options can be used with any datagram socket type in the Internet family.

At the socket level, the socket option `SO_DONTROUTE` may be applied. This option forces datagrams being sent to bypass routing and forwarding by forcing the IP Time To Live field to 1, meaning that the packet will not be forwarded by routers.

Raw IP datagrams can also be sent and received using the TLI connectionless primitives.

Datagrams flow through the IP layer in two directions: from the network *up* to user processes and from user processes *down* to the network. Using this orientation, IP is layered *above* the network interface drivers and *below* the transport protocols such as UDP and TCP. The Internet Control Message Protocol (ICMP) is logically a part of IP. See [icmp\(7P\)](#).

IP provides for a checksum of the header part, but not the data part, of the datagram. The checksum value is computed and set in the process of sending datagrams and checked when receiving datagrams.

IP options in received datagrams are processed in the IP layer according to the protocol specification. Currently recognized IP options include: security, loose source and record route (LSRR), strict source and record route (SSRR), record route, and internet timestamp.

By default, the IP layer will not forward IPv4 packets that are not addressed to it. This behavior can be overridden by using `routeadm(1M)` to enable the `ipv4-forwarding` option. IPv4 forwarding is configured at boot time based on the setting of `routeadm(1M)`'s `ipv4-forwarding` option.

For backwards compatibility, IPv4 forwarding can be enabled or disabled using `ndd(1M)`'s `ip_forwarding` variable. It is set to 1 if IPv4 forwarding is enabled, or 0 if it is disabled.

Additionally, finer-grained forwarding can be configured in IP. Each interface can be configured to forward IP packets by setting the `IFF_ROUTER` interface flag. This flag can be set and cleared using `ifconfig(1M)`'s `router` and `router` options. If an interface's `IFF_ROUTER` flag is set, packets can be forwarded to or from the interface. If it is clear, packets will neither be forwarded from this interface to others, nor forwarded to this interface. Setting the `ip_forwarding` variable sets all of the IPv4 interfaces' `IFF_ROUTER` flags.

For backwards compatibility, each interface creates an `<ifname>:ip_forwarding/dev/ip` variable that can be modified using `ndd(1M)`. An interface's `:ip_forwarding` `ndd` variable is a boolean variable that mirrors the status of its `IFF_ROUTER` interface flag. It is set to 1 if the flag is set, or 0 if it is clear. This interface specific `<ifname> :ip_forwarding` `ndd` variable is obsolete and may be removed in a future release of Solaris. The `ifconfig(1M)` `router` and `-router` interfaces are preferred.

ip(7P)

The IP layer sends an ICMP message back to the source host in many cases when it receives a datagram that can not be handled. A “time exceeded” ICMP message is sent if the “time to live” field in the IP header drops to zero in the process of forwarding a datagram. A “destination unreachable” message is sent if a datagram can not be forwarded because there is no route to the final destination, or if it can not be fragmented. If the datagram is addressed to the local host but is destined for a protocol that is not supported or a port that is not in use, a destination unreachable message is also sent. The IP layer may send an ICMP “source quench” message if it is receiving datagrams too quickly. ICMP messages are only sent for the first fragment of a fragmented datagram and are never returned in response to errors in other ICMP messages.

The IP layer supports fragmentation and reassembly. Datagrams are fragmented on output if the datagram is larger than the maximum transmission unit (MTU) of the network interface. Fragments of received datagrams are dropped from the reassembly queues if the complete datagram is not reconstructed within a short time period.

Errors in sending discovered at the network interface driver layer are passed by IP back up to the user process.

Multi-Data Transmit allows more than one packet to be sent from the IP module to another in a given call, thereby reducing the per-packet processing costs. The behavior of Multi-Data Transmit can be overridden by using `ndd(1M)` to set the `/dev/ip` variable, `ip_multidata_outbound` to 0. Note, the IP module will only initiate Multi-Data Transmit if the network interface driver supports it.

SEE ALSO

`ifconfig(1M)`, `routed(1M)`, `ndd(1M)`, `read(2)`, `write(2)`, `bind(3SOCKET)`, `connect(3SOCKET)`, `getsockopt(3SOCKET)`, `recv(3SOCKET)`, `send(3SOCKET)`, `defaultrouter(4)`, `icmp(7P)`, `if_tcp(7P)`, `inet(7P)`, `ip6(7P)`, `ipsec(7P)`, `routing(7P)`, `tcp(7P)`, `udp(7P)`

Braden, R., *RFC 1122, Requirements for Internet Hosts – Communication Layers*, Information Sciences Institute, University of Southern California, October 1989.

Postel, J., *RFC 791, Internet Protocol – DARPA Internet Program Protocol Specification*, Information Sciences Institute, University of Southern California, September 1981.

DIAGNOSTICS

A socket operation may fail with one of the following errors returned:

EACCES	A <code>bind()</code> operation was attempted with a “reserved” port number and the effective user ID of the process was not the privileged user.
EADDRINUSE	A <code>bind()</code> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
EADDRNOTAVAIL	A <code>bind()</code> operation was attempted for an address that is not configured on this machine.

EINVAL	A <code>sendmsg()</code> operation with a non-NULL <code>msg_accrights</code> was attempted.
EINVAL	A <code>getsockopt()</code> or <code>setsockopt()</code> operation with an unknown socket option name was given.
EINVAL	A <code>getsockopt()</code> or <code>setsockopt()</code> operation was attempted with the IP option field improperly formed; an option field was shorter than the minimum value or longer than the option buffer provided.
EISCONN	A <code>connect()</code> operation was attempted on a socket on which a <code>connect()</code> operation had already been performed, and the socket could not be successfully disconnected before making the new connection.
EISCONN	A <code>sendto()</code> or <code>sendmsg()</code> operation specifying an address to which the message should be sent was attempted on a socket on which a <code>connect()</code> operation had already been performed.
EMSGSIZE	A <code>send()</code> , <code>sendto()</code> , or <code>sendmsg()</code> operation was attempted to send a datagram that was too large for an interface, but was not allowed to be fragmented (such as broadcasts).
ENETUNREACH	An attempt was made to establish a connection by means of <code>connect()</code> , or to send a datagram by means of <code>sendto()</code> or <code>sendmsg()</code> , where there was no matching entry in the routing table; or if an ICMP "destination unreachable" message was received.
ENOTCONN	A <code>send()</code> or <code>write()</code> operation, or a <code>sendto()</code> or <code>sendmsg()</code> operation not specifying an address to which the message should be sent, was attempted on a socket on which a <code>connect()</code> operation had not already been performed.
ENOBUFS	The system ran out of memory for fragmentation buffers or other internal data structures.
ENOBUFS	<code>SO_SNDBUF</code> or <code>SO_RCVBUF</code> exceeds a system limit.
EINVAL	Invalid length for <code>IP_OPTIONS</code> .
EHOSTUNREACH	Invalid address for <code>IP_MULTICAST_IF</code> .
EINVAL	Not a multicast address for <code>IP_ADD_MEMBERSHIP</code> and <code>IP_DROP_MEMBERSHIP</code> .
EADDRNOTAVAIL	Bad interface address for <code>IP_ADD_MEMBERSHIP</code> and <code>IP_DROP_MEMBERSHIP</code> .
EADDRINUSE	Address already joined for <code>IP_ADD_MEMBERSHIP</code> .

ip(7P)

ENOENT	Address not joined for IP_DROP_MEMBERSHIP.
EPERM	No permissions.

NOTES Raw sockets should receive ICMP error packets relating to the protocol; currently such packets are simply discarded.

Users of higher-level protocols such as TCP and UDP should be able to see received IP options.

NAME	ipgpc – IP Generic Packet Classifier
DESCRIPTION	<p>The IP Generic Packet Classifier (<code>ipgpc</code>) module provides packet classification at the Solaris IP layer. <code>ipgpc</code> is an implementation of the Multi-Field (MF) classifier as described in <i>RFC2475: An Architecture for Differentiated Services</i>.</p> <p>The classifier is configured, at startup or dynamically, using a set of “filters.” Filters describe selectors that are matched against input packets that are processed by the classifier. Some selectors feature exact matching data points, while others utilize non-exact or wildcard data points.</p> <p>Each filter is associated with a class describing the next actions to process a packet. There is a many-to-one (M-to-1) mapping relationship between filters and a class. Additionally, each class is aware of which filters are associated with it. A class is configured with a class name and a next action.</p> <p>Unlike traditional classifiers used in edge routers, <code>ipgpc</code> is designed for a host or server device. A host-based classifier provides access to more resources and data than edge routers. User, project, and interface information are available at the host.</p>
STATISTICS	<p>The <code>ipgpc</code> module exports global and per-class statistics (available through <code>kstat</code>):</p> <p>Global statistics:</p> <pre> module: ipgpc instance:<action id> name: ipgpc global stats class: <action name> crtime snaptime nbytes <number of classified bytes> nclasses <number of classes> nfilters <number of filters> npackets <number of classified packets> epackets <number of packets in error> </pre> <p>Per-class statistics:</p> <pre> module: ipgpc instance:<action id> name: <class name> class: <action name> crtime snaptime last match <time of last match> nbytes <number of classified bytes> npackets <number of classified packets> </pre>
FILES	<pre> /kernel/ipp/sparcv9/ipgpc 64-bit module (SPARC only) </pre>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWqos

ipgpc(7ipp)

SEE ALSO [ipqosconf\(1M\)](#), [dlcosmk\(7IPP\)](#), [dscpmk\(7IPP\)](#), [flowacct\(7IPP\)](#), [ipqos\(7IPP\)](#), [tokenmt\(7IPP\)](#), [tswtclmt\(7IPP\)](#)

RFC 2475, An Architecture for Differentiated Services S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss — The Internet Society, 1998

NAME	ipqos – IP Quality of Service						
DESCRIPTION	<p>ipqos is an implementation of the Differentiated Services model defined in <i>RFC2475: An Architecture for Differentiated Services</i>, which defines the following entities: multi-field classifier, meter, marker, and dropper. The Solaris implementation of ipqos adds a flow accounting entity.</p> <p>These entities can be combined into processing paths that constitute a series of actions that are performed on groups of flows. The classifier groups together flows and directs them in a given processing path. Classifier configuration and path construction are achieved using the <code>ipqosconf(1M)</code> command.</p> <p>A summary of the ipqos entities follows. For more information, refer to the corresponding man page for each entity.</p> <p>ipgpc An implementation of the classifier defined in the model. ipgpc has been extended and is able to select traffic based on IP header parameters, user id, project id, interface name, interface group and direction.</p> <p>tokenmt, tswtclmt These modules implement different metering algorithms. tokenmt implements both <i>RFC2697: A Single Rate Three Color Marker</i> and <i>RFC 2698: A Two Rate Three Color Marker</i>. tswtclmt implements <i>RFC2859: A Time Sliding Window Three Color Marker</i>. These modules only implement the metering functions defined in the RFCs.</p> <p>dlcosmk A marker entity that allows the setting of the user priority field of Ethernet frames as defined in the <i>IEEE 802.1D</i> specification. dlcosmk is only available with VLAN capable network interfaces.</p> <p>dscpmk A marker entity that enables the setting of the Differentiated Services Code Point Value in the IP header as defined in <i>RFC 2474: Definition of the Differentiated Services Field (DS Field)</i> in the IPv4 and IPv6 headers.</p> <p>flowacct An accounting module that utilizes the Solaris extended accounting facility. flowacct logs all flows with parameters used to build a charge back mechanism.</p> <p>STATISTICS ipqos modules export statistics through the <code>kstat</code> facility. Exported statistics contain the following common parameters:</p> <table border="0" data-bbox="438 1501 1442 1686"> <tr> <td data-bbox="438 1501 779 1533">module</td> <td data-bbox="779 1501 1442 1533">module name</td> </tr> <tr> <td data-bbox="438 1543 779 1575">instance</td> <td data-bbox="779 1543 1442 1575">dynamic parameter identifying a specific instance</td> </tr> <tr> <td data-bbox="438 1585 779 1617">name</td> <td data-bbox="779 1585 1442 1686">a string for global statistics (for example, ipgpc global stat) or a class name for per-class statistics for a classifier action</td> </tr> </table>	module	module name	instance	dynamic parameter identifying a specific instance	name	a string for global statistics (for example, ipgpc global stat) or a class name for per-class statistics for a classifier action
module	module name						
instance	dynamic parameter identifying a specific instance						
name	a string for global statistics (for example, ipgpc global stat) or a class name for per-class statistics for a classifier action						

ipqos(7ipp)

To verify classifier configuration, generate traffic for each of the configured classes and check that the statistic counters for the expected class are increased. If you're unsure about the parameters for your traffic, you can use `snoop(1M)` to determine them.

Some actions have the instance id of the next configured action in their statistics. This instance id can be used to follow the action processing path. Instance id's -1 and -2 are the built-in actions continue and drop, respectively.

Examples:

To retrieve all statistics for `ipgpc`:

```
kstat -m ipgpc
```

To retrieve statistics for the class `http`:

```
kstat -m ipgpc -c http
```

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWqos (32-bit) SUNWqosx (64-bit)

SEE ALSO

`ipqosconf(1M)`, `dlcosmk(7IPP)`, `dscpmk(7IPP)`, `flowacct(7IPP)`, `ipgpc(7IPP)`, `tokenmt(7IPP)`, `tswtclmt(7IPP)`

RFC 2475, An Architecture for Differentiated Services S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss — The Internet Society, 1998

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers K. Nichols, S. Blake, F. Baker, D. Black — The Internet Society, 1998

RFC 2697, A Single Rate Three Color Marker J. Heinanen, R. Guerin — The Internet Society, 1999

RFC 2698, A Two Rate Three Color Marker J. Heinanen, R. Guerin — The Internet Society, 1999

RFC 2859, A Time Sliding Window Three Colour Marker (TSWTM) W. Fang, N. Seddigh, B. Nandy — The Internet Society, 2000

NAME	iprb – Intel 82557, 82558, 82559–controlled network interface controllers								
SYNOPSIS	<code>/dev/iprb</code>								
DESCRIPTION	The <code>iprb</code> Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code> , over Intel D100 82557, 82558, and 82559 controllers. Multiple 82557, 82558, and 82559 controllers installed within the system are supported by the driver. The <code>iprb</code> driver provides basic support for the 82557, 82558, and 82559 hardware. Functions include chip initialization, frame transmit and receive, multicast support, and error recovery and reporting.								
APPLICATION PROGRAMMING INTERFACE	The cloning, character-special device <code>/dev/iprb</code> is used to access all 82557, 82558, and 82559 devices installed within the system.								
iprb and DLPI	The <code>iprb</code> driver is dependent on <code>/kernel/misc/gld</code> , a loadable kernel module that provides the <code>iprb</code> driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on the primitives supported by the driver. The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> from the user are as follows: <ul style="list-style-type: none"> ■ Maximum SDU is 1500 (ETHERMTU). ■ Minimum SDU is 0. The driver will pad to the mandatory 60-octet minimum packet size. ■ The <code>dlsap</code> address length is 8. ■ MAC type is <code>DL_ETHER</code>. ■ The <code>sap</code> length value is -2, meaning the physical address component is followed immediately by a 2-byte <code>sap</code> component within the DLSAP address. ■ Broadcast address value is Ethernet/IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>). 								
KNOWN PROBLEMS AND LIMITATIONS	x86 based systems with the Intel EtherExpress PRO/100B or the Intel EtherExpress PRO/100+ might hang when the interface is brought down at the very instant that a packet is being received. To avoid this, wait until the system is experiencing light or no network traffic before bringing the interface down. Early versions of the firmware on Intel EtherExpress PRO/100+ and Intel PRO/100+ Management adapters do not support PXE network boot on Solaris systems. Upgrade the firmware if the version is lower than 078. PXE firmware versions are expressed as three-digit build numbers. The build number is typically displayed by the firmware during boot. If the PXE build number is not displayed during boot, change the system BIOS or adapter BIOS configuration to display PXE messages during boot.								
FILES	<table border="0" style="width: 100%;"> <tr> <td style="width: 50%;"><code>iprb</code></td> <td>Device special file</td> </tr> <tr> <td><code>/kernel/drv/iprb.conf</code></td> <td><code>iprb</code> configuration file</td> </tr> <tr> <td><code><sys/stropts.h></code></td> <td><code>stropts</code> network header file</td> </tr> <tr> <td><code><sys/ethernet.h></code></td> <td>Ethernet network header file</td> </tr> </table>	<code>iprb</code>	Device special file	<code>/kernel/drv/iprb.conf</code>	<code>iprb</code> configuration file	<code><sys/stropts.h></code>	<code>stropts</code> network header file	<code><sys/ethernet.h></code>	Ethernet network header file
<code>iprb</code>	Device special file								
<code>/kernel/drv/iprb.conf</code>	<code>iprb</code> configuration file								
<code><sys/stropts.h></code>	<code>stropts</code> network header file								
<code><sys/ethernet.h></code>	Ethernet network header file								

iprb(7D)

<sys/dlpi.h> dlpi network header file

<sys/gld.h> gld network header file

The `iprb.conf` configuration file options include:

-TxURRetry
Default: 3

Allowed Values: 0, 1, 2, 3

Sets the number of retransmissions. Modified when tuning performance.

-MWIEnable
Default: 0 (Disable)

Allowed Values: 0 (Disable), 1 (Enable)

Should only be set for 82558 adapters and systems in which the PCI bus supports Memory Write & Invalidate operations. Can improve the performance for some configurations.

-FlowControl
Default: 0 (Disable)

Allowed Values: 0 (Disable), 1 (Enable)

Setting this value can improve the performance for some configurations

-CollisionBackOffModification
Default: 0 (Disable)

Allowed Values: 0 (Disable), 1 (Enable)

Setting this value can improve the performance for some configurations

-PhyErrataFrequency
Default: 0 (Disable)

Allowed Values: 0 (Disable), 10 (Enable)

If you have problems establishing links with cables length = 70 Ft, set this field to 10

-CpuCycleSaver
Default: 0

Allowed Values: 1 through *FFFFh*

Reasonable Values: *200h* through *800h*

The CPUSaver algorithm improves the system's P/E ratio by reducing the number of interrupts generated by the card. The algorithm bundles multiple receive frames together, then generates a single interrupt for the bundle. Because the microcode

iprb(7D)

does not support run-time configuration, configuration must be done prior to the micro code being loaded into the chip. Changing this value from its default means that the driver will have to be unloaded and loaded for the change to take affect. Setting the CpuCycleSaver option to 0 prevents the algorithm from being used. Because it varies for different network environments, the optimal value for this parameter is impossible to predict. Accordingly, developers should run tests to determine the effect that changing this value has on bandwidth and CPU utilization.

-ForceSpeedDuplex

Default: 5 (Auto-negotiate)

Allowed Values: 4 (100 FDX)

3 (100 HDX)

2 (10 FDX)

1 (10 HDX)

Specify the speed and duplex mode for each instance.

Example: ForceSpeedDuplex=5,4;

Sets iprb0 to autonegotiate and iprb1 to 100 FDX.

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO attributes(5), dlpi(7P), gld(7D)

ipsec(7P)

NAME	ipsec – Internet Protocol Security Architecture
DESCRIPTION	<p>The IP Security Architecture (IPsec) provides protection for IP datagrams. The protection can include confidentiality, strong integrity of the data, partial sequence integrity (replay protection), and data authentication. IPsec is performed inside the IP processing, and it can be applied with or without the knowledge of an Internet application.</p> <p>IPsec applies to both IPv4 and IPv6. See ip(7P) and ip6(7P).</p>
Protection Mechanisms	<p>IPsec provides two mechanisms for protecting data. The Authentication Header (AH) provides strong integrity, replay protection, and data authentication. AH protects as much of the IP datagram as it can. AH cannot protect fields that change nondeterministically between sender and receiver.</p> <p>The Encapsulating Security Payload (ESP) provides confidentiality over what it encapsulates, as well as the services that AH provides, but only over that which it encapsulates. ESP's authentication services are optional, which allow ESP and AH to be used together on the same datagram without redundancy.</p> <p>Authentication and encryption algorithms are used for IPsec. Authentication algorithms produce an integrity checksum value or "digest" based on the data and a key. Encryption algorithms operate on data in units of a "block size."</p>
Security Associations	<p>AH and ESP use Security Associations (SA). SA's are entities that specify security properties from one host to another. Two communicating machines require two SAs (at a minimum) to communicate securely. However, communicating machines that use multicast can share the same multicast SA. SAs are managed through the pf_key(7P) interface. For IPv4, automatic SA management is available through the Internet Key Exchange (IKE), as implemented by in.iked(1M). A command-line front-end is available by means of ipseckey(1M). An IPsec SA is identified by a tuple of <AH or ESP, destination IP address, and SPI>. The Security Parameters Index (SPI) is an arbitrary 32-bit value that is transmitted on the wire with an AH or ESP packet. See ipsecah(7P) or ipsecesp(7P) for an explanation about where the SPI falls in a protected packet.</p>
Protection Policy and Enforcement Mechanisms	<p>Mechanism and policy are separate. The policy for applying IPsec is enforced on a system-wide or per-socket level. Configuring systemwide policy is done via the ipseccnf(1M) command. Configuring per-socket policy is discussed later in this section.</p> <p>Systemwide IPsec policy is applied to incoming and outgoing datagrams. Some additional rules can be applied to outgoing datagrams because of the additional data known by the system. Inbound datagrams can be accepted or dropped. The decision to drop or accept an inbound datagram is based on several criteria which sometimes overlap or conflict. Conflict resolution is resolved by which rule is parsed first, with one exception: if a policy entry states that traffic should bypass all other policy, it is automatically be accepted. Outbound datagrams are sent with or without protection. Protection may (or may not) indicate specific algorithms. If policy normally would protect a datagram, it can be bypassed either by an exception in systemwide policy or by requesting a bypass in per-socket policy.</p>

Intra-machine traffic policies are enforced, but actual security mechanisms are not applied; rather, the outbound policy on an intra-machine packet translates into an inbound packet with those mechanisms applied.

IPsec policy is enforced in the `ip(7P)` driver; several `ndd` tunables for `/dev/ip` affect policy enforcement. These include:

`icmp_accept_clear_messages`

If equal to 1 (the default), allow certain cleartext icmp messages to bypass policy. For ICMP echo requests ("ping" messages), protect the response like the request. If zero, treat icmp messages like other IP traffic.

`igmp_accept_clear_messages`

If 1, allow inbound cleartext IGMP messages to bypass IPsec policy.

`pim_accept_clear_messages`

If 1, allow inbound cleartext PIM messages to bypass IPsec policy.

Per-Socket Policy

The `IP_SEC_OPT` or `IPV6_SEC_OPT` socket option is used to set per-socket IPsec policy. The structure used for an `IP_SEC_OPT` request is:

```
typedef struct ipsec_req {
    uint_t      ipsr_ah_req;          /* AH request */
    uint_t      ipsr_esp_req;        /* ESP request */
    uint_t      ipsr_self_encap_req; /* Self-Encap request */
    uint8_t     ipsr_auth_alg;       /* Auth algs for AH */
    uint8_t     ipsr_esp_alg;        /* Encr algs for ESP */
    uint8_t     ipsr_esp_auth_alg;   /* Auth algs for ESP */
} ipsec_req_t;
```

The IPsec request has fields for both AH and ESP. Algorithms may or may not be specified. The actual request for AH or ESP services can take one of the following values:

`IPSEC_PREF_NEVER`

Bypass all policy. Only the superuser may request this service.

`IPSEC_PREF_REQUIRED`

Regardless of other policy, require the use of the IPsec service.

The following value can be logically ORed to an `IPSEC_PREF_REQUIRED` value:

`IPSEC_PREF_UNIQUE`

Regardless of other policy, enforce a unique SA for traffic originating from this socket.

In the event IP options not normally encapsulated by ESP need to be, the `ipsec_self_encap_req` is used to add an additional IP header outside the original one. Algorithm values from `<net/pfkeyv2.h>` are as follows:

`SADB_AALG_MD5HMAC`

Uses the MD5-HMAC (RFC 2403) algorithm for authentication.

`SADB_AALG_SHA1HMAC`

Uses the SHA1-HMAC (RFC 2404) algorithm for authentication.

ipsec(7P)

SADB_EALG_DESCBC

Uses the DES (*RFC 2405*) algorithm for encryption.

SADB_EALG_3DESCBC

Uses the Triple DES (*RFC 2451*) algorithm for encryption.

SADB_EALG_BLOWFISH

Uses the Blowfish (*RFC 2451*) algorithm for encryption.

SADB_EALG_AES

Uses the Advanced Encryption Standard algorithm for encryption.

An application should use either the `getsockopt(3SOCKET)` or the `setsockopt(3SOCKET)` call to manipulate IPsec requests. For example:

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <net/pfkeyv2.h> /* For SADB_*ALG_* */
/* ... socket setup skipped */
rc = setsockopt(s, IPPROTO_IP, IP_SEC_OPT,
               (const char *)&ipsec_req, sizeof (ipsec_req_t));
```

SECURITY CONSIDERATIONS

While IPsec is an effective tool in securing network traffic, it will not make security problems disappear. Security issues beyond the mechanisms that IPsec offers may be discussed in similar “Security Consideration” sections within individual reference manual pages.

While a non-root user cannot bypass IPsec, a non-root user can set policy to be different from the system-wide policy. For ways to prevent this, consult the `ndd(1M)` variables in `/dev/ip`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`in.iked(1M)`, `ipseconf(1M)`, `ipseckey(1M)`, `ndd(1M)`, `getsockopt(3SOCKET)`, `setsockopt(3SOCKET)`, `attributes(5)`, `inet(7P)`, `ip(7P)`, `ip6(7P)`, `ipsecah(7P)`, `ipsecesp(7P)`, `pf_key(7P)`

Kent, S., and Atkinson, R., *RFC 2401, Security Architecture for the Internet Protocol*, The Internet Society, 1998.

Kent, S. and Atkinson, R., *RFC 2406, IP Encapsulating Security Payload (ESP)*, The Internet Society, 1998.

Madson, C., and Doraswamy, N., *RFC 2405, The ESP DES-CBC Cipher Algorithm with Explicit IV*, The Internet Society, 1998.

Madsen, C. and Glenn, R., *RFC 2403, The Use of HMAC-MD5-96 within ESP and AH*, The Internet Society, 1998.

Madsen, C. and Glenn, R., *RFC 2404, The Use of HMAC-SHA-1-96 within ESP and AH*, The Internet Society, 1998.

Pereira, R. and Adams, R., *RFC 2451, The ESP CBC-Mode Cipher Algorithms*, The Internet Society, 1998.

ipsec(7P)

NAME	ipsec, AH – IPsec Authentication Header						
SYNOPSIS	drv/ipsec						
DESCRIPTION	<p>The ipsec module (AH) provides strong integrity, authentication, and partial sequence integrity (replay protection) to IP datagrams. AH protects the parts of the IP datagram that can be predicted by the sender as it will be received by the receiver. For example, the IP TTL field is not a predictable field, and is not protected by AH.</p> <p>AH is inserted between the IP header and the transport header. The transport header can be TCP, UDP, ICMP, or another IP header, if tunnels are being used. See tun(7M).</p>						
AH Device	AH is implemented as a module that is auto-pushed on top of IP. The entry /dev/ipsec is used for tuning AH with nnd(1M) .						
Authentication Algorithms	Current authentication algorithms supported include HMAC-MD5 and HMAC-SHA-1. Each authentication algorithm has its own key size and key format properties. You can obtain a list of authentication algorithms and their properties by using the ipsecalgs(1M) command. You can also use the functions described in the getipsecalgbyname(3NSL) man page to retrieve the properties of algorithms.						
Security Considerations	Without replay protection enabled, AH is vulnerable to replay attacks. AH does not protect against eavesdropping. Data protected with AH can still be seen by an adversary.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsr</td></tr><tr><td>Interface Stability</td><td>Evolving</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsr	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsr						
Interface Stability	Evolving						
SEE ALSO	ipsecalgs(1M) , ipsecconf(1M) , nnd(1M) , attributes(5) , getipsecalgbyname(3NSL) , tun(7M) , ip(7P) , ipsec(7P) , ipsecesp(7P) Kent, S. and Atkinson, R. <i>RFC 2402, IP Authentication Header</i> , The Internet Society, 1998.						

NAME	ipsecesp, ESP – IPsec Encapsulating Security Payload						
SYNOPSIS	drv/ipsecesp						
DESCRIPTION	<p>The ipsecesp module provides confidentiality, integrity, authentication, and partial sequence integrity (replay protection) to IP datagrams. The encapsulating security payload (ESP) encapsulates its data, enabling it to protect data that follows in the datagram. For TCP packets, ESP encapsulates the TCP header and its data only. If the packet is an IP in IP datagram, ESP protects the inner IP datagram. Per-socket policy allows "self-encapsulation" so ESP can encapsulate IP options when necessary. See ipsec(7P).</p> <p>Unlike the authentication header (AH), ESP allows multiple varieties of datagram protection. (Using a single datagram protection form can expose vulnerabilities.) For example, only ESP can be used to provide confidentiality. But protecting confidentiality alone exposes vulnerabilities in both replay attacks and cut-and-paste attacks. Similarly, if ESP protects only integrity and does not fully protect against eavesdropping, it may provide weaker protection than AH. See ipsecah(7P).</p>						
ESP Device	ESP is implemented as a module that is auto-pushed on top of IP. Use the /dev/ipsecesp entry to tune ESP with ndd(1M).						
Algorithms	ESP uses encryption and authentication algorithms. Authentication algorithms include HMAC-MD5 and HMAC-SHA-1. Encryption algorithms include DES, Triple-DES, Blowfish and AES. Each authentication and encryption algorithm contain key size and key format properties. You can obtain a list of authentication and encryption algorithms and their properties by using the ipsecalgs(1M) command. You can also use the functions described in the getipsecalgbyname(3NSL) man page to retrieve the properties of algorithms. Because of export laws in the United States, not all encryption algorithms are available outside of the United States.						
Security Considerations	ESP without authentication exposes vulnerabilities to cut-and-paste cryptographic attacks as well as eavesdropping attacks. Like AH, ESP is vulnerable to eavesdropping when used without confidentiality.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Availability</td> <td>SUNWcsr (32-bit)</td> </tr> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsr (32-bit)	Interface Stability	Evolving
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsr (32-bit)						
Interface Stability	Evolving						
SEE ALSO	<p>ipsecalgs(1M), ipseconf(1M), ndd(1M), attributes(5), getipsecalgbyname(3NSL), ip(7P), ipsec(7P), ipsecah(7P)</p> <p>Kent, S. and Atkinson, R. <i>RFC 2406, IP Encapsulating Security Payload (ESP)</i>, The Internet Society, 1998.</p>						

isdnio(7I)

NAME	isdnio – ISDN interfaces
SYNOPSIS	<pre>#include <sun/audioio.h> #include <sun/isdnio.h> int ioctl(int fd, int command, /* arg */ ...);</pre>
DESCRIPTION	<p>ISDN <code>ioctl</code> commands are a subset of <code>ioctl(2)</code> commands that perform a variety of control functions on Integrated Services Digital Network (ISDN) STREAMS devices. The arguments <code>command</code> and <code>arg</code> are passed to the file designated by <code>fd</code> and are interpreted by the ISDN device driver.</p> <p><code>fd</code> is an open file descriptor that refers to a stream. <code>command</code> determines the control function to be performed as described in the IOCTLS section of this document. <code>arg</code> represents additional information that is needed by <code>command</code>. The type of <code>arg</code> depends upon the <code>command</code>, but generally it is an integer or a pointer to a command-specific data structure.</p> <p>Since these ISDN commands are a subset of <code>ioctl</code> and <code>streamio(7I)</code>, they are subject to errors as described in those interface descriptions.</p> <p>This set of generic ISDN <code>ioctl</code> commands is meant to control various types of ISDN STREAMS device drivers. The following paragraphs give some background on various types of ISDN hardware interfaces and data formats, and other device characteristics.</p>
Controllers, Interfaces, and Channels	<p>This manual page discusses operations on, and facilities provided by ISDN controllers, interfaces and channels. A controller is usually a hardware peripheral device that provides one or more ISDN interfaces and zero or more auxiliary interfaces. In this context, the term interface is synonymous with the term “port”. Each interface can provide one or more channels.</p>
Time Division Multiplexed Serial Interfaces	<p>ISDN BRI-TE, BRI-NT, and PRI interfaces are all examples of Time Division Multiplexed Serial Interfaces. As an example, a Basic Rate ISDN (BRI) Terminal Equipment (TE) interface provides one D-channel and two B-channels on the same set of signal wires. The BRI interface, at the S reference point, operates at a bit rate of 192,000 bits per second. The bits are encoded using a pseudoternary coding system that encodes a logic one as zero volts, and a logic zero as a positive or negative voltage. Encoding rules state that adjacent logic zeros must be encoded with opposite voltages. Violations of this rule are used to indicate framing information such that there are 4000 frames per second, each containing 48 bits. These 48 bits are divided into channels. Not including framing and synchronization bits, the frame is divided into 8 bits for the B1-channel, 1 bit for the D-channel, 8 bits for B2, 1 bit for D, 8 bits for B1, 1 bit for D, and 8 bits for B2. This results in a 64,000 bps B1-channel, a 64,000 bps B2-channel, and a 16,000 bps D-channel, all on the same serial interface.</p>
Basic Rate ISDN	<p>A Basic Rate ISDN (BRI) interface consists of a 16000 bit per second Delta Channel (D-channel) for signaling and X.25 packet transmission, and two 64000 bit per second Bearer Channels (B-channels) for transmission of voice or data.</p>

The CCITT recommendations on ISDN Basic Rate interfaces, I.430, identify several "reference points" for standardization. From (Stallings89): Reference point T (terminal) corresponds to a minimal ISDN network termination at the customer's premises. It separates the network provider's equipment from the user's equipment. Reference point S (system) corresponds to the interface of individual ISDN terminals. It separates user terminal equipment from network-related communications functions. Reference point R (rate) provides a non-ISDN interface between user equipment that is not ISDN-compatible and adaptor equipment. . . . The final reference point . . . is reference point U (user). This interface describes the full-duplex data signal on the subscriber line.

Some older technology components of some ISDN networks occasionally steal the low order bit of an ISDN B-channel octet in order to transmit in-band signaling information between switches or other components of the network. Even when out-of-band signaling has been implemented in these networks, and the in-band signaling is no longer needed, the bit-robbing mechanism may still be present. This bit robbing behavior does not appreciably affect a voice call, but it will limit the usable bandwidth of a data call to 56000 bits per second instead of 64000 bits per second. These older network components only seem to exist in the United States of America, Canada and Japan. ISDN B-channel data calls that have one end point in the United States, Canada or Japan may be limited to 56000 bps usable bandwidth instead of the normal 64000 bps. Sometimes the ISDN service provider may be able to supply 56kbps for some calls and 64kbps for other calls. On an international call, the local ISDN service provider may advertise the call as 64kbps even though only 56kbps are reliably delivered because of bit-robbing in the foreign ISDN that is not reported to the local switch.

A Basic Rate Interface implements either a Terminal Equipment (TE) interface or a Network Termination (NT) interface. TE's can be ISDN telephones, a Group 4 fax, or other ISDN terminal equipment. A TE connects to an NT in order to gain access to a public or private ISDN network. A private ISDN network, such as provided by a Private Branch Exchange (PBX), usually provides access to the public network.

If multi-point configurations are allowed by an NT, it may be possible to connect up to eight TE's to a single NT interface. All of the TE's in a multipoint configuration share the same D and B-channels. Contention for B-Channels by multiple TEs is resolved by the ISDN switch (NT) through signaling protocols on the D-channel.

Contention for access to the D-channel is managed by a collision detection and priority mechanism. D-channel call control messages have higher priority than other packets. This media access function is managed at the physical layer.

A BRI-TE interface may implement a "Q-channel", the Q-channel is a slow speed, 800 bps, data path from a TE to an NT. Although the structure of the Q-channel is defined in the I.430 specification, the use of the Q-channel is for further study.

A BRI-NT interface may implement an "S-channel", the S-channel is a slow speed, 4000 bps, data path from a NT to an TE. The use of the S-channel is for further study.

isdnio(7I)

Primary Rate ISDN	Primary Rate ISDN (PRI) interfaces are either 1.544Mbps (T1 rate) or 2.048Mbps (E1 rate) and are typically organized as 23 B-channels and one D-Channel (23B+D) for T1 rates, and 30 B-Channels and one D-Channel (30B+D) for E1 rates. The D-channels on a PRI interface operate at 64000 bits per second. T1 rate PRI interface is the standard in the United States, Canada and Japan while E1 rate PRI interface is the standard in European countries. Some E1 rate PRI interface implementations allow access to channel zero which is used for framing.
Channel Types	<p>ISDN channels fall into several categories; D-channels, bearer channels, and management pseudo channels. Each channel has a corresponding device name somewhere under the directory <code>/dev/isdn/</code> as documented in the appropriate hardware specific manual page.</p> <p>D-channels</p> <p>There is at most one D-channel per ISDN interface. The D-channel carries signaling information for the management of ISDN calls and can also carry X.25 packet data. In the case of a PRI interface, there may actually be no D-channel if Non-Facility Associated Signaling is used. D-channels carry data packets that are framed and checked for transmission errors according to the LAP-D protocol. LAP-D uses framing and error checking identical to the High Speed Data Link (HDLC) protocol.</p> <p>B-channels</p> <p>BRI interfaces have two B-channels, B1 and B2. On a BRI interface, the only other type of channel is an H-channel which is a concatenation of the B1 and B2 channels. An H-channel is accessed by opening the "base" channel, B1 in this case, and using the <code>ISDN_SET_FORMAT</code> ioctl to change the configuration of the B-channel from 8-bit, 8 kHz to 16-bit, 8kHz.</p> <p>On a primary rate interface, B channels are numbered from 0 to 31 in Europe and 1 to 23 in the United States, Canada and Japan.</p> <p>H-Channels</p> <p>A BRI or PRI interface can offer multiple B-channels concatenated into a single, higher bandwidth channel. These concatenated B-channels are referred to as an "H-channels" on a BRI interface. The PRI interface version of an H-channel is referred to as an Hn-channels where n is a number indicating how the B-channels have been aggregated into a single channel.</p> <ul style="list-style-type: none">■ A PRI interface H0 channel is 384 kbps allowing 3H0+D on a T1 rate PRI interface and 4H0+D channels on an E1 rate PRI interface.■ A T1 PRI interface H11 channel is 1536 kbps (24×64000bps). This will consume the channel normally reserved for the D-channel, so signaling must be done with Non-Facility Associated Signaling (NFAS) from another PRI interface.■ An E1 PRI interface H12 channel is 1920 kbps (30×64000bps). An H12-channel leaves room for the framing-channel as well as the D-channel. <p>Auxiliary channels</p> <p>Auxiliary channels are non-ISDN hardware interfaces that are closely tied to the ISDN interfaces. An example would be a video or audio coder/decoder (codec). The existence of an auxiliary channel usually implies that one or more B-channels can be "connected" to an auxiliary interface in hardware.</p>

Management pseudo-channels

A management pseudo-channel is used for the management of a controller, interface, or hardware channel. Management channels allow for out-of-band control of hardware interfaces and for out-of-band notification of status changes. There is at least one management device per hardware interface.

There are three different types of management channels implemented by ISDN hardware drivers:

- A controller management device handles all ioctls that simultaneously affect hardware channels on different interfaces. Examples include resetting a controller, mu-code (as in the Greek letter mu) downloading of a controller, or the connection of an ISDN B-channel to an auxiliary channel that represents an audio coder/decoder (codec). The latter case would be accomplished using the `ISDN_SET_CHANNEL` ioctl.
- An interface management device handles all ioctls that affect multiple channels on the same interface. Messages associated with the activation and deactivation of an interface arrive on the management device associated with the D channel of an ISDN interface.
- Auxiliary interfaces may also have management devices. See the hardware specific man pages for operations on auxiliary devices.

Trace pseudo-channels

A device driver may choose to implement a trace device for a data or management channel. Trace channels receive a special `M_PROTO` header with the original channel's original `M_PROTO` or `M_DATA` message appended to the special header. The header is described by:

```
typedef struct {
    uint_t    seq;      /* Sequence number */
    int      type;     /* device dependent */
    struct    timeval  timestamp;
    char     _f[8];    /* filler */
} audtrace_hdr_t;
```

ISDN Channel types

The `isdn_chan_t` type enumerates the channels available on ISDN interfaces. If a particular controller implements any auxiliary channels then those auxiliary channels will be described in a controller specific manual page. The defined channels are described by the `isdn_chan_t` type as shown below:

```
/* ISDN channels */
typedef enum {
    ISDN_CHAN_NONE = 0x0, /* No channel given */
    ISDN_CHAN_SELF, /* The channel performing the ioctl */
    ISDN_CHAN_HOST, /* Unix STREAM */
    ISDN_CHAN_CTRL_MGT, /* Controller management */

    /* TE channel defines */

    ISDN_CHAN_TE_MGT, /* Receives activation/deactivation */
    ISDN_CHAN_TE_D_TRACE, /* Trace device for protocol analysis apps */
    ISDN_CHAN_TE_D,
    ISDN_CHAN_TE_B1,
```

isdnio(7I)

```
ISDN_CHAN_TE_B2,

/* NT channel defines */

ISDN_CHAN_NT_MGT,      /* Receives activation/deactivation */
ISDN_CHAN_NT_D_TRACE, /* Trace device for protocol analysis apps */
ISDN_CHAN_NT_D,
ISDN_CHAN_NT_B1,
ISDN_CHAN_NT_B2,

/* Primary rate ISDN */

ISDN_CHAN_PRI_MGT,
ISDN_CHAN_PRI_D,
ISDN_CHAN_PRI_B0, ISDN_CHAN_PRI_B1,
ISDN_CHAN_PRI_B2, ISDN_CHAN_PRI_B3,
ISDN_CHAN_PRI_B4, ISDN_CHAN_PRI_B5,
ISDN_CHAN_PRI_B6, ISDN_CHAN_PRI_B7,
ISDN_CHAN_PRI_B8, ISDN_CHAN_PRI_B9,
ISDN_CHAN_PRI_B10, ISDN_CHAN_PRI_B11,
ISDN_CHAN_PRI_B12, ISDN_CHAN_PRI_B13,
ISDN_CHAN_PRI_B14, ISDN_CHAN_PRI_B15,
ISDN_CHAN_PRI_B16, ISDN_CHAN_PRI_B17,
ISDN_CHAN_PRI_B18, ISDN_CHAN_PRI_B19,
ISDN_CHAN_PRI_B20, ISDN_CHAN_PRI_B21,
ISDN_CHAN_PRI_B22, ISDN_CHAN_PRI_B23,
ISDN_CHAN_PRI_B24, ISDN_CHAN_PRI_B25,
ISDN_CHAN_PRI_B26, ISDN_CHAN_PRI_B27,
ISDN_CHAN_PRI_B28, ISDN_CHAN_PRI_B29,
ISDN_CHAN_PRI_B30, ISDN_CHAN_PRI_B31,

/* Auxiliary channel defines */

ISDN_CHAN_AUX0, ISDN_CHAN_AUX1, ISDN_CHAN_AUX2, ISDN_CHAN_AUX3,
ISDN_CHAN_AUX4, ISDN_CHAN_AUX5, ISDN_CHAN_AUX6, ISDN_CHAN_AUX7
} isdn_chan_t;
```

ISDN Interface types

The `isdn_interface_t` type enumerates the interfaces available on ISDN controllers. The defined interfaces are described by the `isdn_interface_t` type as shown below:

```
/* ISDN interfaces */
typedef enum {
    ISDN_TYPE_UNKNOWN = -1, /* Not known or applicable */
    ISDN_TYPE_SELF = 0, /*
        * For queries, application may
        * put this value into "type" to
        * query the state of the file
        * descriptor used in an ioctl.
        */
    ISDN_TYPE_OTHER, /* Not an ISDN interface */
    ISDN_TYPE_TE,
    ISDN_TYPE_NT,
    ISDN_TYPE_PRI,
} isdn_interface_t;
```

Activation and Deactivation of ISDN Interfaces

The management device associated with an ISDN D-channel is used to request activation, deactivation and receive information about the activation state of the interface. See the descriptions of the `ISDN_PH_ACTIVATE_REQ` and

ISDN_MPH_DEACTIVATE_REQ ioctls. Changes in the activation state of an interface are communicated to the D-channel application through M_PROTO messages sent up-stream on the management device associated with the D-channel. If the D-channel protocol stack is implemented as a user process, the user process can retrieve the M_PROTO messages using the `getmsg(2)` system call.

These M_PROTO messages have the following format:

```
typedef struct isdn_message {
    unsigned int magic;           /* set to ISDN_PROTO_MAGIC */
    isdn_interface_t type;       /* Interface type */
    isdn_message_type_t message; /* CCITT or vendor Primitive */
    unsigned int vendor[5];      /* Vendor specific content */
} isdn_message_t;
typedef enum isdn_message_type {
    ISDN_VPH_VENDOR = 0, /* Vendor specific messages */
    ISDN_PH_AI,          /* Physical: Activation Ind */
    ISDN_PH_DI,          /* Physical: Deactivation Ind */
    ISDN_MPH_AI,         /* Management: Activation Ind */
    ISDN_MPH_DI,         /* Management: Deactivation Ind */
    ISDN_MPH_EI1,        /* Management: Error 1 Indication */
    ISDN_MPH_EI2,        /* Management: Error 2 Indication */
    ISDN_MPH_II_C,       /* Management: Info Ind, connection */
    ISDN_MPH_II_D,       /* Management: Info Ind, disconn. */
} isdn_message_type_t;
```

IOCTLS

STREAMS IOCTLS

All of the [streamio\(7I\)](#) ioctl commands may be issued for a device conforming to the the isdnio interface.

ISDN interfaces that allow access to audio data should implement a reasonable subset of the [audio\(7I\)](#) interface.

ISDN ioctls

ISDN_PH_ACTIVATE_REQ

Request ISDN physical layer activation. This command is valid for both TE and NT interfaces. *fd* must be a D-channel file descriptor. *arg* is ignored.

TE activation will occur without use of the ISDN_PH_ACTIVATE_REQ ioctl if the device corresponding to the TE D-channel is open, "on", and the ISDN switch is requesting activation.

ISDN_MPH_DEACTIVATE_REQ

fd must be an NT D-channel file descriptor. *arg* is ignored.

This command requests ISDN physical layer de-activation. This is not valid for TE interfaces. A TE interace may be turned off by use of the ISDN_PARAM_POWER command or by `close(2)` on the associated *fd*.

ISDN_ACTIVATION_STATUS

fd is the file descriptor for a D-channel, the management device associated with an ISDN interface, or the management device associated with the controller. *arg* is a pointer to an `isdn_activation_status_t` structure. Although it is possible for applications to determine the current activation state with this ioctl, a D-channel protocol stack should instead process messages from the management pseudo channel associated with the D-channel.

```
typedef struct isdn_activation_status {
    isdn_interface_t type;
    enum isdn_activation_state activation;
} isdn_activation_status_t;
typedef enum isdn_activation_state {
    ISDN_OFF = 0,          /* Interface is powered down */
    ISDN_UNPLUGGED,      /* Power but no-physical connection */
    ISDN_DEACTIVATED_REQ, /* Pending Deactivation, NT Only */
    ISDN_DEACTIVATED,    /* Activation is permitted */
    ISDN_ACTIVATE_REQ,   /* Attempting to activate */
    ISDN_ACTIVATED,      /* Interface is activated */
} isdn_activation_state_t;
```

The `type` field should be set to `ISDN_TYPE_SELF`. The device specific interface type will be returned in the `type` field.

The `isdn_activation_status_t` structure contains the interface type and the current activation state. `type` is the interface type and should be set by the caller to `ISDN_TYPE_SELF`.

ISDN_INTERFACE_STATUS

The `ISDN_INTERFACE_STATUS` ioctl retrieves the status and statistics of an ISDN interface. The requesting channel must own the interface whose status is being requested or the ioctl will fail. *fd* is the file descriptor for an ISDN interface management device. *arg* is a pointer to a `struct isdn_interface_info`. If the `interface` field is set to `ISDN_TYPE_SELF`, it will be changed in the returned structure to reflect the proper device-specific interface of the requesting *fd*.

```
typedef struct isdn_interface_info {
    isdn_interface_t interface;
    enum isdn_activation_state activation;
    unsigned int ph_ai; /* Physical: Activation Ind */
    unsigned int ph_di; /* Physical: Deactivation Ind */
    unsigned int mph_ai; /* Management: Activation Ind */
    unsigned int mph_di; /* Management: Deactivation Ind */
    unsigned int mph_ei1; /* Management: Error 1 Indication */
    unsigned int mph_ei2; /* Management: Error 2 Indication */
    unsigned int mph_ii_c; /* Management: Info Ind, connection */
    unsigned int mph_ii_d; /* Management: Info Ind, disconn. */
} isdn_interface_info_t;
```

ISDN_CHANNEL_STATUS

The `ISDN_CHANNEL_STATUS` ioctl retrieves the status and statistics of an ISDN channel. The requesting channel must own the channel whose status is being

requested or the ioctl will fail. *fd* is any file descriptor. *arg* is a pointer to a struct `isdn_channel_info`. If the interface field is set to `ISDN_CHAN_SELF`, it will be changed in the returned structure to reflect the proper device-specific channel of the requesting *fd*.

```
typedef struct isdn_channel_info {
    isdn_chan_t    channel;
    enum isdn_iostate    iostate;
    struct isdn_io_stats {
        ulong_t    packets; /* packets transmitted or received */
        ulong_t    octets; /* octets transmitted or received */
        ulong_t    errors; /* errors packets transmitted or received */
    } transmit, receive;
} isdn_channel_info_t;
```

ISDN_PARAM_SET

fd is the file descriptor for a management device. *arg* is a pointer to a struct `isdn_param`. This command allows the setting of various ISDN physical layer parameters such as timers. This command uses the same arguments as the `ISDN_PARAM_GET` command.

ISDN_PARAM_GET

fd is the file descriptor for a management device. *arg* is a pointer to a struct `isdn_param`. This command provides for querying the value of a particular ISDN physical layer parameter.

```
typedef enum {
    ISDN_PARAM_NONE = 0,
    ISDN_PARAM_NT_T101, /* NT Timer, 5-30 s, in milliseconds */
    ISDN_PARAM_NT_T102, /* NT Timer, 25-100 ms, in milliseconds */
    ISDN_PARAM_TE_T103, /* TE Timer, 5-30 s, in milliseconds */
    ISDN_PARAM_TE_T104, /* TE Timer, 500-1000 ms, in milliseconds */
    ISDN_PARAM_MAINT, /* Manage the TE Maintenance Channel */
    ISDN_PARAM_ASMB, /* Modify Activation State Machine Behavior */
    ISDN_PARAM_POWER, /* Take the interface online or offline */
    ISDN_PARAM_PAUSE, /* Paused if == 1, else not paused == 0 */
} isdn_param_tag_t;
enum isdn_param_asmb {
    ISDN_PARAM_TE_ASMB_CCITT88, /* 1988 bluebook */
    ISDN_PARAM_TE_ASMB_CTS2, /* Conformance Test Suite 2 */
};
typedef struct isdn_param {
    isdn_param_tag_t    tag;
    union {
        unsigned int us; /* micro seconds */
        unsigned int ms; /* Timer value in ms */
        unsigned int flag; /* Boolean */
    };
    enum isdn_param_asmb    asmb;
    enum isdn_param_maint    maint;
    struct {
        isdn_chan_t    channel; /* Channel to Pause */
        int    paused; /* TRUE or FALSE */
    } pause;
    unsigned int reserved[2]; /* reserved, set to zero */
} value;
```

isdnio(7I)

```
} isdn_param_t;
```

ISDN_PARAM_POWER

If an implementation provides power on and off functions, then power should be on by default. If `flag` is `ISDN_PARAM_POWER_OFF` then a TE interface is forced into state F0, NT interfaces are forced into state G0. If `flag` is `ISDN_PARAM_POWER_ON` then a TE interface will immediately transition to state F3 when the TE D-channel is opened. If `flag` is one, an NT interface will transition to state G1 when the NT D-channel is opened.

Implementations that do not provide `ISDN_POWER` return failure with `errno` set to `ENXIO`. `ISDN_POWER` is different from `ISDN_PH_ACTIVATE_REQ` since CCITT specification requires that if a BRI-TE interface device has power, then it permits activation.

ISDN_PARAM_NT_T101

This parameter accesses the NT timer value T1. The CCITT recommendations specify that timer T1 has a value from 5 to 30 seconds. Other standards may differ.

ISDN_PARAM_NT_T102

This parameter accesses the NT timer value T2. The CCITT recommendations specify that timer T2 has a value from 25 to 100 milliseconds. Other standards may differ.

ISDN_PARAM_TE_T103

This parameter accesses the TE timer value T3. The CCITT recommendations specify that timer T3 has a value from 5 to 30 seconds. Other standards may differ.

ISDN_PARAM_TE_T104

This parameter accesses the TE timer value T4. The CTS2 specifies that timer T4 is either not used or has a value from 500 to 1000 milliseconds. Other standards may differ. CTS2 requires that timer T309 be implemented if T4 is not available.

ISDN_PARAM_MAINT

This parameter sets the multi-framing mode of a BRI-TE interface. For normal operation this parameter should be set to `ISDN_PARAM_MAINT_ECHO`. Other uses of this parameter are dependent on the definition and use of the BRI interface S and Q channels.

ISDN_PARAM_ASMB

There are a few differences in the BRI-TE interface activation state machine standards. This parameter allows the selection of the appropriate standard. At this time, only `ISDN_PARAM_TE_ASMB_CCITT88` and `ISDN_PARAM_TE_ASMB_CTS2` are available.

ISDN_PARAM_PAUSE

This parameter allows a management device to pause the IO on a B-channel. `pause.channel` is set to indicate which channel is to be paused or un-paused. `pause.paused` is set to zero to un-pause and one to pause. `fd` is associated with an ISDN interface management device. `arg` is a pointer to a `struct isdn_param`.

ISDN_SET_LOOPBACK

fd is the file descriptor for an ISDN interface's management device. *arg* is a pointer to an `isdn_loopback_request_t` structure.

```
typedef enum {
    ISDN_LOOPBACK_LOCAL,
    ISDN_LOOPBACK_REMOTE,
} isdn_loopback_type_t;
typedef enum {
    ISDN_LOOPBACK_B1 = 0x1,
    ISDN_LOOPBACK_B2 = 0x2,
    ISDN_LOOPBACK_D = 0x4,
    ISDN_LOOPBACK_E_ZERO = 0x8,
    ISDN_LOOPBACK_S = 0x10,
    ISDN_LOOPBACK_Q = 0x20,
} isdn_loopback_chan_t;
typedef struct isdn_loopback_request {
    isdn_loopback_type_t type;
    int channels;
} isdn_loopback_request_t;
```

An application can receive D-channel data during D-Channel loopback but cannot transmit data. The field type is the bitwise OR of at least one of the following values:

```
ISDN_LOOPBACK_B1    (0x1)  /* loopback on B1-channel */
ISDN_LOOPBACK_B2    (0x2)  /* loopback on B2-channel */
ISDN_LOOPBACK_D     (0x4)  /* loopback on D-channel */
ISDN_LOOPBACK_E_ZERO (0x8)  /* force E-channel to Zero if */
                        /* fd is for NT interface */
ISDN_LOOPBACK_S     (0x10) /* loopback on S-channel */
ISDN_LOOPBACK_Q     (0x20) /* loopback on Q-channel */
```

ISDN_RESET_LOOPBACK

arg is a pointer to an `isdn_loopback_request_t` structure.

`ISDN_RESET_LOOPBACK` turns off the selected loopback modes.

ISDN Data Format

The `isdn_format_t` type is meant to be a complete description of the various data modes and rates available on an ISDN interface. Several macros are available for setting the format fields. The `isdn_format_t` structure is shown below:

```
/* ISDN channel data format */
typedef enum {
    ISDN_MODE_NOTSPEC, /* Not specified */
    ISDN_MODE_HDLC,    /* HDLC framing and error checking */
    ISDN_MODE_TRANSPARENT /* Transparent mode */
} isdn_mode_t;

/* Audio encoding types (from audioio.h) */

#define AUDIO_ENCODING_NONE (0) /* no encoding*/
#define AUDIO_ENCODING_ULAW (1) /* mu-law */
#define AUDIO_ENCODING_ALAW (2) /* A-law */
#define AUDIO_ENCODING_LINEAR (3) /* Linear PCM */
```

isdnio(7I)

```
typedef struct isdn_format {
    isdn_mode_t    mode;
    unsigned int   sample_rate; /* sample frames/sec*/
    unsigned int   channels;    /* # interleaved chans */
    unsigned int   precision;   /* bits per sample */
    unsigned int   encoding;    /* data encoding */
} isdn_format_t;
/*
 * These macros set the fields pointed
 * to by the macro argument (isdn_format_t*)fp in preparation
 * for the ISDN_SET_FORMAT ioctl.
 */
ISDN_SET_FORMAT_BRI_D(fp)      /* BRI D-channel */
ISDN_SET_FORMAT_PRI_D(fp)      /* PRI D-channel */
ISDN_SET_FORMAT_HDL_C_B64(fp)  /* BRI B-ch @ 56kbps */
ISDN_SET_FORMAT_HDL_C_B56(fp)  /* BRI B-ch @ 64kbps */
ISDN_SET_FORMAT_VOIC_E_ULAW(fp) /* BRI B-ch voice */
ISDN_SET_FORMAT_VOIC_E_ALAW(fp) /* BRI B-ch voice */
ISDN_SET_FORMAT_BRI_H(fp)      /* BRI H-channel */
```

ISDN Datapath Types

Every STREAMS stream that carries data to or from the ISDN serial interfaces is classified as a channel-stream datapath. A possible ISDN channel-stream datapath device name for a TE could be `/dev/isdn/0/te/b1`.

On some hardware implementations, it is possible to route the data from hardware channel to hardware channel completely within the chip or controller. This is classified as a channel-channel datapath. There does not need to be any open file descriptor for either channel in this configuration. Only when data enters the host and utilizes a STREAMS stream is this classified as an ISDN channel-stream datapath.

ISDN Management Stream

A management stream is a STREAMS stream that exists solely for control purposes and is not intended to carry data to or from the ISDN serial interfaces. A possible management device name for a TE could be `/dev/isdn/0/te/mgt`.

Channel Management IOCTLS

The following ioctls describe operations on individual channels and the connection of multiple channels.

ISDN_SET_FORMAT

fd is a data channel, the management pseudo-channel associated with the data channel, or the management channel associated with the data channel's interface or controller. *arg* is a pointer to a `struct isdn_format_req`. The ISDN_SET_FORMAT ioctl sets the format of an ISDN channel-stream datapath. It may be issued on both an open ISDN channel-stream datapath Stream or an ISDN Management Stream. Note that an `open(2)` call for a channel-stream datapath will fail if an ISDN_SET_FORMAT has never been issued after a reset, as the mode for all channel-stream datapaths is initially biased to ISDN_MODE_NOTSPEC. *arg* is a pointer to an ISDN format type (`isdn_format_req_t`).

```
typedef struct isdn_format_req {
    isdn_chan_t    channel;
    isdn_format_t  format; /* data format */
    int reserved[4]; /* future use - must be 0 */
} isdn_format_req_t;
```

If there is not an open channel-stream datapath for a requested channel, the default format of that channel will be set for a subsequent `open(2)`.

To modify the format of an open `STREAM`, the driver will disconnect the hardware channel, flush the internal hardware queues, set the new default configuration, and finally reconnect the data path using the newly specified format. Upon taking effect, all state information will be reset to initial conditions, as if a channel was just opened. It is suggested that the user flush the interface as well as consult the hardware specific documentation to insure data integrity.

If a user desires to connect more than one B channel, such as an H-channel, the B-channel with the smallest offset should be specified, then the precision should be specified multiples of 8. For an H-channel the precision value would be 16. The user should subsequently open the base B-channel. If any of the sequential B-channels are busy the open will fail, otherwise all of the B-channels that are to be used in conjunction will be marked as busy.

The returned failure codes and their descriptions are listed below:

```

EPERM  /* No permission for intended operation */
EINVAL /* Invalid format request */
EIO    /* Set format attempt failed. */

```

ISDN_SET_CHANNEL

The `ISDN_SET_CHANNEL` `ioctl` sets up a data connection within an ISDN controller. The `ISDN_SET_CHANNEL` `ioctl` can only be issued from an ISDN management stream to establish or modify channel-channel datapaths. The `ioctl` parameter `arg` is a pointer to an ISDN connection request (`isdn_conn_req_t*`). Once a data path is established, data flow is started as soon as the path endpoints become active. Upon taking effect, all state information is reset to initial conditions, as if a channel was just opened.

The `isdn_conn_req_t` structure is shown below. The five fields include the receive and transmit ISDN channels, the number of directions of the data path, as well as the data format. The reserved field must always be set to zero.

```

/* Number of directions for data flow */
typedef enum {
    ISDN_PATH_NOCHANGE = 0, /* Invalid value */
    ISDN_PATH_DISCONNECT, /* Disconnect data path */
    ISDN_PATH_ONEWAY,     /* One way data path */
    ISDN_PATH_TWOWAY,     /* Bi-directional data path */
} isdn_path_t;
typedef struct isdn_conn_req {
    isdn_chan_t from;
    isdn_chan_t to;
    isdn_path_t dir; /* uni/bi-directional or disconnect */
    isdn_format_t format; /* data format */
    int reserved[4]; /* future use - must be 0 */
} isdn_conn_req_t;

```

isdnio(7I)

To specify a read-only, write-only, or read-write path, or to disconnect a path, the `dir` field should be set to `ISDN_PATH_ONEWAY`, `ISDN_PATH_TWOWAY`, and `ISDN_PATH_DISCONNECT` respectively. To modify the format of a channel-channel datapath, a user must disconnect the channel and then reconnect with the desired format.

The returned failure codes and their descriptions are listed below:

```
EPERM /* No permission for intended operation */
EBUSY /* Connection in use */
EINVAL /* Invalid connection request */
EIO /* Connection attempt failed */
```

ISDN_GET_FORMAT

The `ISDN_GET_FORMAT` ioctl gets the ISDN data format of the channel-stream datapath described by `fd`. `arg` is a pointer to an ISDN data format request type (`isdn_format_req_t*`). `ISDN_GET_FORMAT` can be issued on any channel to retrieve the format of any channel it owns. For example, if issued on the TE management channel, the format of any other te channel can be retrieved.

ISDN_GETCONFIG

The `ISDN_GETCONFIG` ioctl is used to get the current connection status of all ISDN channels associated with a particular management STREAM. `ISDN_GETCONFIG` also retrieves a hardware identifier and the generic interface type. `arg` is an ISDN connection table pointer (`isdn_conn_tab_t*`). The `isdn_conn_tab_t` structure is shown below:

```
typedef struct isdn_conn_tab {
    char name[ISDN_ID_SIZE]; /* identification string */
    isdn_interface_t type;
    int maxpaths; /* size in entries of app's array int npaths; */
    /* number of valid entries returned by driver */
    isdn_conn_req_t *paths; /* connection table in app's memory */
} isdn_conn_tab_t;
```

The table contains a string which is the interface's unique identification string. The second element of this table contains the ISDN transmit and receive connections and configuration for all possible data paths for each type of ISDN controller hardware. Entries that are not connected will have a value of `ISDN_NO_CHAN` in the `from` and `to` fields. The number of entries will always be `ISDN_MAX_CHANS`, and can be referenced in the hardware specific implementation documentation. An `isdn_conn_tab_t` structure is allocated on a per controller basis.

SEE ALSO [getmsg\(2\)](#), [ioctl\(2\)](#), [open\(2\)](#), [poll\(2\)](#), [read\(2\)](#), [write\(2\)](#), [audio\(7I\)](#), [dbri\(7D\)](#), [streamio\(7I\)](#)

ISDN, An Introduction – William Stallings, Macmillan Publishing Company. ISBN 0-02-415471-7

NAME	isp – ISP SCSI Host Bus Adapter Driver
SYNOPSIS	
Sbus	QLGC,isp@ <i>sbus-slot</i> ,10000
PCI	SUNW,isp _{two} @ <i>pci-slot</i>
DESCRIPTION	<p>The ISP Host Bus Adapter is a SCSI compliant nexus driver that supports the Qlogic ISP1000 SCSI and the ISP1040B SCSI chips. The ISP1000 chip works on SBus and the ISP1040B chip works on PCI bus. The ISP is an intelligent SCSI Host Bus Adapter chip that reduces the amount of CPU overhead used in a SCSI transfer.</p> <p>The <code>isp</code> driver supports the standard functions provided by the SCSI interface. The driver supports tagged and untagged queuing, fast and wide SCSI, and auto request sense, but does not support linked commands. The PCI version ISP Host bus adapter based on ISP1040B also supports Fast-20 SCSI devices.</p>
CONFIGURATION	<p>The <code>isp</code> driver can be configured by defining properties in <code>isp.conf</code> which override the global SCSI settings. Supported properties are <code>scsi-options</code>, <code>target<n>-scsi-options</code>, <code>scsi-reset-delay</code>, <code>scsi-watchdog-tick</code>, <code>scsi-tag-age-limit</code>, <code>scsi-initiator-id</code>, and <code>scsi-selection-timeout</code>.</p> <p><code>target<n>-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target<n></code>. <code><n></code> is a hex value that can vary from 0 to f. Refer to <code>scsi_hba_attach(9F)</code> for details.</p> <p>Both the ISP1000 and ISP1040B support only certain SCSI selection timeout values. The valid values are 25, 50, 75, 100, 250, 500, 750 and 1000. These properties are in units of milliseconds.</p>
EXAMPLES	<p>EXAMPLE 1 SCSI Options</p> <p>Create a file called <code>/kernel/drv/isp.conf</code> and add this line:</p> <pre>scsi-options=0x78;</pre> <p>This will disable tagged queuing, fast SCSI, and Wide mode for all <code>isp</code> instances. The following will disable an option for one specific ISP (refer to <code>driver.conf(4)</code>):</p> <pre>name="isp" parent="/iommu@f,e0000000/sbus@f,e0001000" reg=1,0x10000,0x450 target1-scsi-options=0x58 scsi-options=0x178 scsi-initiator-id=6;</pre> <p>Note that the default initiator ID in OBP is 7 and that the change to ID 6 will occur at attach time. It may be preferable to change the initiator ID in OBP.</p> <p>The above would set <code>scsi-options</code> for target 1 to 0x58 and for all other targets on this SCSI bus to 0x178.</p> <p>The physical pathname of the parent can be determined using the <code>/devices</code> tree or following the link of the logical device name:</p>

EXAMPLE 1 SCSI Options *(Continued)*

```
example# ls -l /dev/rdisk/c2t0d0s0
lrwxrwxrwx 1 root root 76 Aug 22 13:29 /dev/rdisk/c2t0d0s0 ->
../../../../devices/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/sd@0,0:a,raw
```

Determine the register property values using the output of `prtconf(1M)` with the `-v` option:

```
QLGC,isp, instance #0
...
Register Specifications:
    Bus Type=0x1, Address=0x10000, Size=450
```

EXAMPLE 2 ISP Properties

The `isp` driver exports properties indicating per target the negotiated transfer speed (`target<n>-sync-speed`), whether tagged queuing has been enabled (`target<n>-TQ`), and whether the wide data transfer has been negotiated (`target<n>-wide`). The `sync-speed` property value is the data transfer rate in KB/sec. The `target-TQ` and `target-wide` properties have no value. The existence of these properties indicate that tagged queuing or wide transfer has been enabled. Refer to `prtconf(1M)` (verbose option) for viewing the `isp` properties.

```
QLGC,isp, instance #2
Driver software properties:
    name <target0-TQ> length <0> -- <no value>.
    name <target0-wide> length <0> -- <no value>.
    name <target0-sync-speed> length <4>
        value <0x000028f5>.
    name <scsi-options> length <4>
        value <0x000003f8>.
    name <scsi-watchdog-tick> length <4>
        value <0x0000000a>.
    name <scsi-tag-age-limit> length <4>
        value <0x00000008>.
    name <scsi-reset-delay> length <4>
        value <0x00000bb8>.
```

EXAMPLE 3 PCI Bus

To achieve the same setting of SCSI-options as in instance #0 above on a PCI machine, create a file called `/kernel/drv/isp.conf` and add the following entries.

```
name="isp" parent="/pci@1f,2000/pci@1"
    unit-address="4"
    scsi-options=0x178
    target3-scsi-options=0x58 scsi-initiator-id=6;
```

The physical pathname of the parent can be determined using the `/devices` tree or following the link of the logical device name:

To set `scsi-options` more specifically per device type, add the following line in the `/kernel/drv/isp.conf` file:

EXAMPLE 3 PCI Bus *(Continued)*

```
device-type-scsi-options-list =
    "SEAGATE ST32550W", "seagate-scsi-options" ;
seagate-scsi-options = 0x58;
```

All device which are of this specific disk type will have `scsi-options` set to 0x58.

`scsi-options` specified per target ID has the highest precedence, followed by `scsi-options` per device type. Global (for all `isp` instances) `scsi-options` per bus has the lowest precedence.

The system needs to be rebooted before the specified `scsi-options` take effect.

EXAMPLE 4 Driver Capabilities

The target driver needs to set capabilities in the `isp` driver in order to enable some driver features. The target driver can query and modify these capabilities: `synchronous`, `tagged-qing`, `wide-xfer`, `auto-rqsense`, `qfull-retries`, `qfull-retry-interval`. All other capabilities can only be queried.

By default, `tagged-qing`, `auto-rqsense`, and `wide-xfer` capabilities are disabled, while `disconnect`, `synchronous`, and `untagged-qing` are enabled. These capabilities can only have binary values (0 or 1). The default values for `qfull-retries` and `qfull-retry-interval` are both 10. The `qfull-retries` capability is a `uchar_t` (0 to 255) while `qfull-retry-interval` is a `ushort_t` (0 to 65535).

The target driver needs to enable `tagged-qing` and `wide-xfer` explicitly. The `untagged-qing` capability is always enabled and its value cannot be modified, because `isp` can queue commands even when `tagged-qing` is disabled.

Whenever there is a conflict between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only `whom != 0` is supported in the `scsi_ifsetcap(9F)` call.

Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

FILES /kernel/drv/isp ELF Kernel Module
/kernel/drv/isp.conf Configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO `prtconf(1M)`, `driver.conf(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_reset(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

isp(7D)

Writing Device Drivers

ANSI Small Computer System Interface-2 (SCSI-2)

QLogic Corporation, ISP1000 Firmware Interface Specification

QLogic Corporation, ISP1020 Firmware Interface Specification

QLogic Corporation, ISP1000 Technical Manual

QLogic Corporation, ISP1020a/1040a Technical Manual

QLogic Corporation, Differences between the ISP1020a/1040a and the ISP1020B/1040B - Application Note

DIAGNOSTICS

The messages described below may appear on the system console as well as being logged.

The first set of messages may be displayed while the `isp` driver is first trying to attach. All of these messages mean that the `isp` driver was unable to attach. These messages are preceded by "`isp<number>`", where "`<number>`" is the instance number of the ISP Host Bus Adapter.

`Device in slave-only slot, unused`

The SBus device has been placed in a slave-only slot and will not be accessible; move to non-slave-only SBus slot.

`Device is using a hilevel intr, unused`

The device was configured with an interrupt level that cannot be used with this `isp` driver. Check the device.

`Failed to alloc soft state`

Driver was unable to allocate space for the internal state structure. Driver did not attach to device; SCSI devices will be inaccessible.

`Bad soft state`

Driver requested an invalid internal state structure. Driver did not attach to device; SCSI devices will be inaccessible.

`Unable to map registers`

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

`Cannot add intr`

Driver was not able to add the interrupt routine to the kernel. Driver did not attach to device; SCSI devices will be inaccessible.

`Unable to attach`

Driver was unable to attach to the hardware for some reason that may be printed. Driver did not attach to device; SCSI devices will be inaccessible.

The next set of messages can be displayed at any time. They will be printed with the full device pathname followed by the shorter form described above.

Firmware should be < 0x<number> bytes
 Firmware size exceeded allocated space and will not download firmware. This could mean that the firmware was corrupted somehow. Check the `isp` driver.

Firmware checksum incorrect
 Firmware has an invalid checksum and will not be downloaded.

Chip reset timeout
 ISP chip failed to reset in the time allocated; may be bad hardware.

Stop firmware failed
 Stopping the firmware failed; may be bad hardware.

Load ram failed
 Unable to download new firmware into the ISP chip.

DMA setup failed
 The DMA setup failed in the host adapter driver on a `scsi_pkt`. This will return `TRAN_BADPKT` to a SCSI target driver.

Bad request pkt
 The ISP Firmware rejected the packet as being set up incorrectly. This will cause the `isp` driver to call the target completion routine with the reason of `CMD_TRAN_ERR` set in the `scsi_pkt`. Check the target driver for correctly setting up the packet.

Bad request pkt header
 The ISP Firmware rejected the packet as being set up incorrectly. This will cause the `isp` driver to call the target completion routine with the reason of `CMD_TRAN_ERR` set in the `scsi_pkt`. Check the target driver for correctly setting up the packet.

Polled command timeout on <number>. <number>
 A polled command experienced a timeout. The target device, as noted by the target lun (<number>.<number>) information, may not be responding correctly to the command, or the ISP chip may be hung. This will cause an error recovery to be initiated in the `isp` driver. This could mean a bad device or cabling.

SCSI Cable/Connection problem

Hardware/Firmware error
 The ISP chip encountered a firmware error of some kind. The problem is probably due to a faulty scsi cable or improper cable connection. This error will cause the `isp` driver to do error recovery by resetting the chip.

Received unexpected SCSI Reset
 The ISP chip received an unexpected SCSI Reset and has initiated its own internal error recovery, which will return all the `scsi_pkt` with reason set to `CMD_RESET`.

Fatal timeout on target <number>. <number>
 The `isp` driver found a command that had not completed in the correct amount of time; this will cause error recovery by the `isp` driver. The device that experienced the timeout was at target lun (<number>.<number>).

Fatal error, resetting interface
 This is an indication that the `isp` driver is doing error recovery. This will cause all outstanding commands that have been transported to the `isp` driver to be

isp(7D)

completed via the `scsi_pkt` completion routine in the target driver with reason of `CMD_RESET` and status of `STAT_BUS_RESET` set in the `scsi_pkt`.

NAME	ixgb – SUNWixgb, 10G Gigabit Ethernet driver for Intel 82597ex controllers and Sun Kirkwood adapters
SYNOPSIS	/dev/ixgb
DESCRIPTION	<p>The ixgb 10Gigabit Ethernet driver is a multi-threaded, loadable, clonable, GLD-based STREAMS driver supporting the Data Link Provider Interface, DLPI, on Intel 82597ex 10 Gigabit Ethernet controllers and Sun 10G Kirkwood adapters. The Intel 10G controller incorporates both MAC and PHY functions and provides 10G (fiber) Ethernet operation on the SR and LR connectors. The Sun Kirkwood adapter is a 133 MHz PCIX 10 Gig Ethernet card utilizing the Intel 82597EX PCIX MAC controller with XFP-based 10GigE optics.</p> <p>The ixgb driver functions include controller initialization, frame transmit and receive, promiscuous and multicast support and error recovery and reporting.</p> <p>The ixgb driver and hardware support auto-negotiation, a protocol specified by the <i>802.ae</i> specification.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device /dev/ixgb is used to access all Intel 10G controllers and Sun Kirkwood adapters installed within the system.</p> <p>The ixgb driver is dependent on /kernel/misc/gld, a loadable kernel module that provides the ixgb driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on supported primitives.</p> <p>You must send an explicit DL_ATTACH_REQ message to associate the opened stream with a particular device (PPA). The PPA ID is interpreted as an unsigned integer data type and indicates the corresponding device instance (unit) number. The driver returns an error (DL_ERROR_ACK) if the PPA field value does not correspond to a valid device instance number for the system. The device is initialized on first attach and de-initialized (stopped) at last detach.</p> <p>The values returned by the driver in the DL_INFO_ACK primitive in response to the DL_INFO_REQ are as follows:</p> <ul style="list-style-type: none"> ■ Maximum SDU is 1500 (ETHERMTU - defined in <sys/ethernet.h>). ■ Minimum SDU is 0. ■ DLSAP address length is 8 . ■ MAC type is DL_ETHER . ■ SAP length value is -2 meaning the physical address component is followed immediately by a 2 byte SAP component within the DLSAP address. ■ Broadcast address value is Ethernet/IEEE broadcast address (FF:FF:FF:FF:FF:FF). <p>Once in the DL_ATTACHED state, you must send a DL_BIND_REQ to associate a particular Service Access Point (SAP) with the stream.</p> <p>By default, the ixgb driver performs auto-negotiation to select the 10G link speed.</p>

ixgb(7d)

FILES /dev/ixgb ixgb special character device.
/kernel/drv/sparcv9/ixgb/ixgb driver binary.
/kernel/drv/amd64/ixgb 64-bit kernel module (x86 only.)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86

SEE ALSO `attributes(5)`, `gld(7D)`, `streamio(7I)`, `dlpi(7P)`

Writing Device Drivers

STREAMS Programming Guide

Network Interfaces Programmer's Guide

NAME	jfb – XVR-1200 Graphics Accelerator device driver	
DESCRIPTION	<p>The <code>jfb</code> module is the device driver for the Sun XVR-600 and Sun XVR-1200 Graphics Accelerators. The XVR-1200 Graphics Accelerator is a high resolution, high performance PCI graphics framebuffer providing hardware 2D and 3D acceleration. Sun XVR-600 is the single pipeline version of the Sun XVR-1200.</p> <p>The <code>jfbdaemon</code> process loads the <code>jfb</code> microcode at system startup time and during the resume sequence of a suspend-resume cycle.</p>	
FILES	<code>/dev/fbs/jfb0</code>	Device special file for XVR-600 and XVR-1200 high performance single screen.
	<code>/dev/fbs/jfb0a</code>	Device special file for the XVR-1200 first video out.
	<code>/dev/fbs/jfb0b</code>	Device special file for the XVR-1200 second video out.
	<code>/usr/lib/jfb.ucode</code>	<code>jfb</code> microcode.
	<code>/usr/sbin/jfbdaemon</code>	<code>jfb</code> microcode loader.
SEE ALSO	SUNWjfb_config(1M)	

jfca(7D)

NAME jfca – JNI Fibre Channel Adapter (FCA) Driver

SYNOPSIS SUNW,jfca@pci-slot

DESCRIPTION The jfca host bus adapter driver is a Sun Fibre Channel transport layer-compliant nexus driver for the JNI FCX-6562-L and FCX2-6562-L adapters. These adapters support Fibre Channel SCSI and IP Protocols, FC-AL public loop profile, point-to-point, and fabric connection and Fibre Channel service classes two and three.

The jfca driver interfaces with the Sun Fibre Channel transport layer to support the standard functions provided by the SCSA interface. It supports auto request sense and tagged queueing by default. The driver requires that all devices have unique hard addresses in private loop configurations.

FILES /kernel/drv/sparcv9/jfca
64-bit ELF kernel module

/kernel/drv/jfca.conf
Configuration file

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC
Availability	SUNWjfca

SEE ALSO prtconf(1M), driver.conf(4), fcp(7D), fp(7D)

Writing Device Drivers

ANSI X3.230:1994, Fibre Channel Physical Signaling (FC-PH)

Project 1134-D, Fibre Channel Generic Services (FC-GS-2)

ANSI X3.269-1996, Fibre Channel Arbitrated Loop (FC-AL)

ANSI X3.270-1996, Fibre Channel Protocol for SCSI (FCP-SCSI)

ANSI X3.270-1996, SCSI-3 Architecture Model (SAM)

Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)

Fabric Loop Attachment (FC-FLA)

NAME	kb – keyboard STREAMS module										
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stream.h> #include <sys/stropts.h> #include <sys/vuid_event.h> #include <sys/kbio.h> #include <sys/kbd.h> ioctl(fd, I_PUSH, "kb");</pre>										
DESCRIPTION	<p>The kb STREAMS module processes byte streams generated by a keyboard attached to a CPU serial port. Definitions for altering keyboard translation and reading events from the keyboard are contained in <code><sys/kbio.h></code> and <code><sys/kbd.h></code>.</p> <p>The kb STREAMS module utilizes a set of keyboard tables to recognize which keys have been typed. Each translation table is an array of 128 16-bit words (unsigned shorts). If a table entry is less than 0x100, the entry is treated as an ISO 8859/1 character. Higher values indicate special characters that invoke more complicated actions.</p>										
Keyboard Translation Mode	<p>The keyboard can be in one of the following translation modes:</p> <p><code>TR_NONE</code> Keyboard translation is turned off and up/down key codes are reported.</p> <p><code>TR_ASCII</code> ISO 8859/1 codes are reported.</p> <p><code>TR_EVENT</code> <code>firm_events</code> are reported.</p> <p><code>TR_UNTRANS_EVENT</code> <code>firm_events</code> containing unencoded keystation codes are reported for all input events within the window system.</p>										
Keyboard Translation-Table Entries	<p>All instances of the kb module share seven translation tables that convert raw keystation codes to event values. The tables are:</p> <table border="0"> <tr> <td>Unshifted</td> <td>Used when a key is depressed and no shifts are in effect.</td> </tr> <tr> <td>Shifted</td> <td>Used when a key is depressed and a Shift key is held down.</td> </tr> <tr> <td>Caps Lock</td> <td>Used when a key is depressed and Caps Lock is in effect.</td> </tr> <tr> <td>Alt Graph</td> <td>Used when a key is depressed and the Alt Graph key is held down.</td> </tr> <tr> <td>Num Lock</td> <td>Used when a key is depressed and Num Lock is in effect.</td> </tr> </table>	Unshifted	Used when a key is depressed and no shifts are in effect.	Shifted	Used when a key is depressed and a Shift key is held down.	Caps Lock	Used when a key is depressed and Caps Lock is in effect.	Alt Graph	Used when a key is depressed and the Alt Graph key is held down.	Num Lock	Used when a key is depressed and Num Lock is in effect.
Unshifted	Used when a key is depressed and no shifts are in effect.										
Shifted	Used when a key is depressed and a Shift key is held down.										
Caps Lock	Used when a key is depressed and Caps Lock is in effect.										
Alt Graph	Used when a key is depressed and the Alt Graph key is held down.										
Num Lock	Used when a key is depressed and Num Lock is in effect.										

kb(7M)

Controlled Used when a key is depressed and the Control key is held down. (Regardless of whether a Shift key or the Alt Graph is being held down, or whether Caps Lock or Num Lock is in effect).

Key Up Used when a key is released.

Each key on the keyboard has a `key station` code that represents a number from 0 to 127. The number is used as an index into the translation table that is currently in effect. If the corresponding entry in the translation table is a value from 0 to 255, the value is treated as an ISO 8859/1 character, and the character is the result of the translation.

If the entry in the translation table is higher than 255, it is a special entry. Special entry values are classified according to the value of the high-order bits. The high-order value for each class is defined as a constant, as shown below. When added to the constant, the value of the low-order bits distinguish between keys within each class:

SHIFTKEYS 0x100

A shift key. The value of the particular shift key is added to determine which shift mask to apply:

CAPSLOCK 0	Caps Lock key.
SHIFTLOCK 1	“Shift Lock” key.
LEFTSHIFT 2	Left-hand Shift key.
RIGHTSHIFT 3	Right-hand Shift key.
LEFTCTRL 4	Left-hand (or only) Control key.
RIGHTCTRL 5	Right-hand Control key.
ALTGRAPH 9	Alt Graph key.
ALT 10	Alternate or Alt key.
NUMLOCK 11	Num Lock key.

BUCKYBITS 0x200

Used to toggle mode-key-up/down status without altering the value of an accompanying ISO 8859/1 character. The actual bit-position value, minus 7, is added.

METABIT 0 The Meta key was pressed along with the key. This is the only user-accessible bucky bit. It is ORed in as the 0x80 bit; since this bit is a legitimate bit in a character, the only way to distinguish between, for example, 0xA0 as META+0x20 and 0xA0 as an 8-bit character is to watch for META key up and META key down events and keep track of whether the META key was down.

SYSTEMBIT 1 The System key was pressed. This is a place holder to indicate which key is the system-abort key.

FUNNY 0x300

Performs various functions depending on the value of the low 4 bits:

NOP 0x300	Does nothing.
OOPS 0x301	Exists, but is undefined.
HOLE 0x302	There is no key in this position on the keyboard, and the position-code should not be used.
RESET 0x306	Keyboard reset.
ERROR 0x307	The keyboard driver detected an internal error.
IDLE 0x308	The keyboard is idle (no keys down).
COMPOSE 0x309	The COMPOSE key; the next two keys should comprise a two-character COMPOSE key sequence.
NONL 0x30A	Used only in the Num Lock table; indicates that this key is not affected by the Num Lock state, so that the translation table to use to translate this key should be the one that would have been used had Num Lock not been in effect.
0x30B — 0x30F	Reserved for non-parameterized functions.

FA_CLASS 0x400

A floating accent or “dead key.” When this key is pressed, the next key generates an event for an accented character; for example, “floating accent grave” followed by the “a” key generates an event with the ISO 8859/1 code for the “a with grave accent” character. The low-order bits indicate which accent; the codes for the individual “floating accents” are as follows:

FA_UMLAUT 0x400	umlaut
FA_CFLEX 0x401	circumflex
FA_TILDE 0x402	tilde
FA_CEDILLA 0x403	cedilla
FA_ACUTE 0x404	acute accent
FA_GRAVE 0x405	grave accent

STRING 0x500

The low-order bits index a table of strings. When a key with a STRING entry is depressed, the characters in the null-terminated string for that key are sent, character-by-character. The maximum length is defined as:

KTAB_STRLEN	10 Individual string numbers are defined as:
HOMEARROW 0x00	
UPARROW 0x01	

kb(7M)

DOWNARROW
0x02

LEFTARROW
0x03

RIGHTARROW
0x04

String numbers 0x05 — 0x0F are available for custom entries.

FUNCKEYS 0x600

There are 64 keys reserved for function keys. The actual positions are usually on the left/right/top/bottom of the keyboard.

The next-to-lowest 4 bits indicate the group of function keys:

LEFTFUNC 0x600

RIGHTFUNC 0x610

TOPFUNC 0x610 0x610

BOTTOMFUNC 0x630 The low 4 bits indicate the function key number within the group:

LF(*n*) (LEFTFUNC+(*n*)-1)

RF(*n*) (RIGHTFUNC+(*n*)-1)

TF(*n*) (TOPFUNC+(*n*)-1)

BF(*n*) (BOTTOMFUNC+(*n*)-1)

PADKEYS 0x700

A "numeric keypad key." These entries should appear only in the Num Lock translation table; when Num Lock is in effect, these events will be generated by pressing keys on the right-hand keypad. The low-order bits indicate which key. The codes for the individual keys are:

PADEQUAL 0x700 "=" key

PADSLASH 0x701 "/" key

PADSTAR 0x702 "*" key

PADMINUS 0x703 "-" key

PADSEP 0x704 "," key

PAD7 0x705 "7" key

PAD8 0x706 "8" key

PAD9 0x707 "9" key

PADPLUS 0x708 "+" key

PAD4 0x709 "4" key

PAD5 0x70A	“5” key
PAD6 0x70B	“6” key
PAD1 0x70C	“1” key
PAD2 0x70D	“2” key
PAD3 0x70E	“3” key
PAD0 0x70F	“0” key
PADDOT 0x710	“.” key
PADENTER 0x711	“Enter” key

When a function key is pressed in `TR_ASCII` mode, the following escape sequence is sent:

```
ESC[0 . . . 9z
```

where `ESC` is a single escape character and “0...9” indicates the decimal representation of the function-key value. For example, function key `R1` sends the sequence:

```
ESC[208z
```

because the decimal value of `RF(1)` is 208. In `TR_EVENT` mode, if there is a `VUID` event code for the function key in question, an event with that event code is generated; otherwise, individual events for the characters of the escape sequence are generated.

Keyboard Compatibility Mode

When started, the `kb STREAMS` module is in the compatibility mode. When the keyboard is in the `TR_EVENT` translation mode, ISO 8859/1 characters from the upper half of the character set (that is, characters with the eighth bit set), are presented as events with codes in the `ISO_FIRST` range (as defined in `<<sys/vuid_event.h>>`). For backwards compatibility with older versions of the keyboard driver, the event code is `ISO_FIRST` plus the character value. When compatibility mode is turned off, ISO 8859/1 characters are presented as events with codes equal to the character code.

DESCRIPTION

The following `ioctl()` requests set and retrieve the current translation mode of a keyboard:

<code>KIOCTRANS</code>	Pointer to an <code>int</code> . The translation mode is set to the value in the <code>int</code> pointed to by the argument.
<code>KIOCGTRANS</code>	Pointer to an <code>int</code> . The current translation mode is stored in the <code>int</code> pointed to by the argument.

`ioctl()` requests for changing and retrieving entries from the keyboard translation table use the `kiockeymap` structure:

```
struct kiockeymap {
int    kio_tablemask; /* Translation table (one of: 0, CAPSMASK,
                      * SHIFTMASK, CTRLMASK, UPMASK,
```

kb(7M)

```
        * ALTGRAPHMASK, NUMLOCKMASK)
        */
#define KIOABORT1 -1      /* Special "mask": abort1 keystation */
#define KIOABORT2 -2      /* Special "mask": abort2 keystation */
        uchar_t kio_station; /* Physical keyboard key station (0-127) */
        ushort_t kio_entry; /* Translation table station's entry */
        char kio_string[10]; /* Value for STRING entries-null terminated */
};
```

KIOCSKEY Pointer to a `kiockeymap` structure. The translation table entry referred to by the values in that structure is changed. The `kio_tablemask` request specifies which of the following translation tables contains the entry to be modified:

UPMASK 0x0080	“Key Up” translation table.
NUMLOCKMASK 0x0800	“Num Lock” translation table.
CTRLMASK 0x0030	“Controlled” translation table.
ALTGRAPHMASK 0x0200	“Alt Graph” translation table.
SHIFTMASK 0x000E	“Shifted” translation table.
CAPSMASK 0x0001	“Caps Lock” translation table.
(No shift keys pressed or locked)	“Unshifted” translation table.

The `kio_station` request specifies the keystation code for the entry to be modified. The value of `kio_entry` is stored in the entry in question. If `kio_entry` is between `STRING` and `STRING+15`, the string contained in `kio_string` is copied to the appropriate string table entry. This call may return `EINVAL` if there are invalid arguments.

Special values of `kio_tablemask` can affect the two step “break to the PROM monitor” sequence. The usual sequence is `L1-a` or `Stop`. If `kio_tablemask` is `KIOABORT1`, then the value of `kio_station` is set to be the first keystation in the sequence. If `kio_tablemask`, is `KIOABORT2` then the value of `kio_station` is set to be the second keystation in the sequence. An attempt to change the “break to the PROM monitor” sequence without having superuser permission results in an `EPERM` error.

KIOCGKEY	The argument is a pointer to a <code>kiockeymap</code> structure. The current value of the keyboard translation table entry specified by <code>kio_tablemask</code> and <code>kio_station</code> is stored in the structure pointed to by the argument. This call may return <code>EINVAL</code> if there are invalid arguments.
KIOCTYPE	<p>The argument is a pointer to an <code>int</code>. A code indicating the type of the keyboard is stored in the <code>int</code> pointed to by the argument:</p> <p><code>KB_SUN3</code> Sun Type 3 keyboard</p> <p><code>KB_SUN4</code> Sun Type 4 or 5 keyboard, or non-USB Sun Type 6 keyboard</p> <p><code>KB_USB</code> USB standard HID keyboard, including Sun Type 6 USB keyboards</p> <p><code>KB_ASCII</code> ASCII terminal masquerading as keyboard</p> <p><code>KB_PC</code> Type 101 PC keyboard</p> <p><code>KB_DEFAULT</code> Stored in the <code>int</code> pointed to by the argument if the keyboard type is unknown. In case of error, -1 is stored in the <code>int</code> pointed to by the argument.</p>
KIOCLAYOUT	The argument is a pointer to an <code>int</code> . On a Sun Type 4 keyboard, the layout code specified by the keyboard's DIP switches is stored in the <code>int</code> pointed to by the argument.
KIOCCMD	<p>The argument is a pointer to an <code>int</code>. The command specified by the value of the <code>int</code> pointed to by the argument is sent to the keyboard. The commands that can be sent are:</p> <p>Commands to the Sun Type 3 and Sun Type 4 keyboards:</p> <p><code>KBD_CMD_RESET</code> Reset keyboard as if power-up.</p> <p><code>KBD_CMD_BELL</code> Turn on the bell.</p> <p><code>KBD_CMD_NOBELL</code> Turn off the bell.</p> <p><code>KBD_CMD_CLICK</code> Turn on the click annunciator.</p> <p><code>KBD_CMD_NOCLICK</code> Turn off the click annunciator.</p> <p>Commands to the Sun Type 4 keyboard:</p> <p><code>KBD_CMD_SETLED</code> Set keyboard LEDs.</p> <p><code>KBD_CMD_GETLAYOUT</code> Request that keyboard indicate layout.</p>

kb(7M)

Inappropriate commands for particular keyboard types are ignored. Since there is no reliable way to get the state of the bell or click (because the keyboard cannot be queried and a process could do writes to the appropriate serial driver — circumventing this `ioctl()` request) an equivalent `ioctl()` to query its state is not provided.

KIOCSLED The argument is a pointer to an `char`. On the Sun Type 4 keyboard, the LEDs are set to the value specified in that `char`. The values for the four LEDs are:

`LED_CAPS_LOCK` “Caps Lock” light.

`LED_COMPOSE` “Compose” light.

`LED_SCROLL_LOCK` “Scroll Lock” light.

`LED_NUM_LOCK` “Num Lock” light.

On some Japanese layouts, the value for the fifth LED is:

`LED_KANA` “Kana” light.

KIOCGLED Pointer to a `char`. The current state of the LEDs is stored in the `char` pointed to by the argument.

KIOCSCOMPAT Pointer to an `int`. “Compatibility mode” is turned on if the `int` has a value of 1, and is turned off if the `int` has a value of 0.

KIOCGCOMPAT Pointer to an `int`. The current state of “compatibility mode” is stored in the `int` pointed to by the argument.

The following `ioctl()` request allows the default effect of the keyboard abort sequence to be changed.

KIOCSKABORTEN

Pointer to an `int`. The keyboard abort sequence effect (typically L1-A or Stop-A on the keyboard on SPARC systems, F1-A on x86 systems, and BREAK on the serial console device) is enabled if the `int` has a value of `KIOCSKABORTENABLE(1)`. If the value is `KIOCSKABORTDISABLE(0)`, the keyboard abort sequence effect is disabled. If the value is `KIOCSKABORTALTERNATE(2)`, the Alternate Break sequence is in effect and is defined by the serial console drivers `zs(7D)`, `se(7D)` and `asy(7D)`. Any other value of the parameter for this `ioctl()` is treated as `enable`. The Alternate Break sequence is applicable to the serial console devices only.

Due to a risk of incorrect sequence interpretation, SLIP and certain other binary protocols should not be run over the serial console port when Alternate Break sequence is in effect. Although PPP is a binary protocol, it is able to avoid these sequences using the ACCM feature in *RFC 1662*. For Solaris PPP 4.0, you do this by adding the following line to the `/etc/ppp/options` file (or other configuration files used for the connection; see `pppd(1M)` for details):

asynomap 0x00002000

SLIP has no comparable capability, and must not be used if the Alternate Break sequence is in use.

This `ioctl()` will be active and retain state even if there is no physical keyboard in the system. The default effect (`enable`) causes the operating system to suspend and enter the kernel debugger (if present) or the system prom (on most systems with OpenBoot proms). The default effect is enabled on most systems, but may be different on server systems with key switches in the 'secure' position. On these systems, the effect is always disabled when the key switch is in the 'secure' position. This `ioctl()` returns `EPERM` if the caller is not the superuser.

These `ioctl()` requests are supported for compatibility with the system keyboard device `/dev/kbd`.

`KIOCSDIRECT`

Has no effect.

`KIOCGDIRECT`

Always returns 1.

The following `ioctl()` requests are used to set and get the keyboard autorepeat delay and rate.

`KIOCSRPTDELAY`

This argument is a pointer to an `int`, which is the kb autorepeat delay, unit in millisecond.

`KIOCGRPTDELAY`

This argument is a pointer to an `int`. The current auto repeat delay setting is stored in the integer pointed by the argument, unit in millisecond.

`KIOCSRPRATE`

This argument is a pointer to an `int`, which is the kb autorepeat rate, unit in millisecond.

`KIOCGRPRATE`

This argument is a pointer to an `int`. The current auto repeat rate setting is stored in the integer pointed by the argument, unit in millisecond.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Stable

SEE ALSO `kbd(1)`, `loadkeys(1)`, `kadb(1M)`, `pppd(1M)`, `keytables(4)`, `attributes(5)`, `zs(7D)`, `se(7D)`, `asy(7D)`, `termio(7I)`, `usbkbm(7M)`

NOTES Many keyboards released after Sun Type 4 keyboard also report themselves as Sun Type 4 keyboards.

kdmouse(7D)

NAME | kdmouse – built-in mouse device interface

DESCRIPTION | The kdmouse driver supports machines with built-in PS/2 mouse interfaces. It allows applications to obtain information about the mouse's movements and the status of its buttons.

Programs are able to read directly from the device. The data returned corresponds to the byte sequences as defined in the *IBM PS/2 Technical Reference Manual*.

FILES | /dev/kdmouse device file

ATTRIBUTES | See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO | `attributes(5)`, `vuidmice(7M)`
IBM PS/2 Technical Reference Manual.

- NAME** kmdb – Kernel debugger
- DESCRIPTION** The kmdb driver is the mechanism used by mdb to invoke and control kmdb. This is *not* a public interface.
- FILES** /dev/kmdb
Kernel debugger driver.
- ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Private

- SEE ALSO** `attributes(5)`

kstat(7D)

NAME	kstat – kernel statistics driver
DESCRIPTION	The <code>kstat</code> driver is the mechanism used by the <code>kstat(3KSTAT)</code> library to extract kernel statistics. This is NOT a public interface.
FILES	<code>/dev/kstat</code> kernel statistics driver
SEE ALSO	<code>kstat(3KSTAT)</code> , <code>kstat(9S)</code>

NAME	ksyms – kernel symbols						
SYNOPSIS	<code>/dev/ksyms</code>						
DESCRIPTION	<p>The file <code>/dev/ksyms</code> is a character special file that allows read-only access to an ELF format image containing two sections: a symbol table and a corresponding string table. The contents of the symbol table reflect the symbol state of the currently running kernel. You can determine the size of the image with the <code>fstat()</code> system call. The recommended method for accessing the <code>/dev/ksyms</code> file is by using the ELF access library. See <code>elf(3ELF)</code> for details. If you are not familiar with ELF format, see <code>a.out(4)</code>.</p> <p><code>/dev/ksyms</code> is an executable for the processor on which you are accessing it. It contains ELF program headers which describe the text and data segment(s) in kernel memory. Since <code>/dev/ksyms</code> has no text or data, the fields specific to file attributes are initialized to <code>NULL</code>. The remaining fields describe the text or data segment(s) in kernel memory.</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">Symbol table</td> <td>The SYMTAB section contains the symbol table entries present in the currently running kernel. This section is ordered as defined by the ELF definition with locally-defined symbols first, followed by globally-defined symbols. Within symbol type, the symbols are ordered by kernel module load time. For example, the kernel file symbols are first, followed by the first module's symbols, and so on, ending with the symbols from the last module loaded.</td> </tr> <tr> <td></td> <td>The section header index (<code>st_shndx</code>) field of each symbol entry in the symbol table is set to <code>SHN_ABS</code>, because any necessary symbol relocations are performed by the kernel link editor at module load time.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">String table</td> <td>The STRTAB section contains the symbol name strings that the symbol table entries reference.</td> </tr> </table>	Symbol table	The SYMTAB section contains the symbol table entries present in the currently running kernel. This section is ordered as defined by the ELF definition with locally-defined symbols first, followed by globally-defined symbols. Within symbol type, the symbols are ordered by kernel module load time. For example, the kernel file symbols are first, followed by the first module's symbols, and so on, ending with the symbols from the last module loaded.		The section header index (<code>st_shndx</code>) field of each symbol entry in the symbol table is set to <code>SHN_ABS</code> , because any necessary symbol relocations are performed by the kernel link editor at module load time.	String table	The STRTAB section contains the symbol name strings that the symbol table entries reference.
Symbol table	The SYMTAB section contains the symbol table entries present in the currently running kernel. This section is ordered as defined by the ELF definition with locally-defined symbols first, followed by globally-defined symbols. Within symbol type, the symbols are ordered by kernel module load time. For example, the kernel file symbols are first, followed by the first module's symbols, and so on, ending with the symbols from the last module loaded.						
	The section header index (<code>st_shndx</code>) field of each symbol entry in the symbol table is set to <code>SHN_ABS</code> , because any necessary symbol relocations are performed by the kernel link editor at module load time.						
String table	The STRTAB section contains the symbol name strings that the symbol table entries reference.						
SEE ALSO	<code>kernel(1M)</code> , <code>stat(2)</code> , <code>elf(3ELF)</code> , <code>kvm_open(3KVM)</code> , <code>a.out(4)</code> , <code>mem(7D)</code>						
WARNINGS	<p>The kernel is dynamically configured. It loads kernel modules when necessary. Because of this aspect of the system, the symbol information present in the running system can vary from time to time, as kernel modules are loaded and unloaded.</p> <p>When you open the <code>/dev/ksyms</code> file, you have access to an ELF image which represents a snapshot of the state of the kernel symbol information at that instant in time. While the <code>/dev/ksyms</code> file remains open, kernel module autounloading is disabled, so that you are protected from the possibility of acquiring stale symbol data. Note that new modules can still be loaded, however. If kernel modules are loaded while you have the <code>/dev/ksyms</code> file open, the snapshot held by you will not be</p>						

ksyms(7D)

updated. In order to have access to the symbol information of the newly loaded modules, you must first close and then reopen the `/dev/ksyms` file. Be aware that the size of the `/dev/ksyms` file will have changed. You will need to use the `fstat()` function (see `stat(2)`) to determine the new size of the file.

Avoid keeping the `/dev/ksyms` file open for extended periods of time, either by using `kvm_open(3KVM)` of the default namelist file or with a direct open. There are two reasons why you should not hold `/dev/ksyms` open. First, the system's ability to dynamically configure itself is partially disabled by the locking down of loaded modules. Second, the snapshot of symbol information held by you will not reflect the symbol information of modules loaded after your initial open of `/dev/ksyms`.

Note that the `ksyms` driver is a loadable module, and that the kernel driver modules are only loaded during an open system call. Thus it is possible to run `stat(2)` on the `/dev/ksyms` file without causing the `ksyms` driver to be loaded. In this case, the file size returned is `UNKNOWN_SIZE`. A solution for this behavior is to first open the `/dev/ksyms` file, causing the `ksyms` driver to be loaded (if necessary). You can then use the file descriptor from this open in a `fstat()` system call to get the file's size.

NOTES The kernel virtual memory access library (`libkvm`) routines use `/dev/ksyms` as the default namelist file. See `kvm_open(3KVM)` for details.

NAME	ldterm – standard STREAMS terminal line discipline module
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/termios.h> int ioctl (fd, I_PUSH, "ldterm");</pre>
DESCRIPTION	<p>The ldterm STREAMS module provides most of the termio(7I) terminal interface. The vis module does not perform the low-level device control functions specified by flags in the <code>c_cflag</code> word of the <code>termio/termios</code> structure, or by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code>, or <code>INPCK</code> flags in the <code>c_iflag</code> word of the <code>termio/termios</code> structure. Those functions must be performed by the driver or by modules pushed below the ldterm module. The ldterm module performs all other <code>termio/termios</code> functions, though some may require the cooperation of the driver or modules pushed below ldterm and may not be performed in some cases. These include the <code>IXOFF</code> flag in the <code>c_iflag</code> word and the delays specified in the <code>c_oflag</code> word.</p> <p>The ldterm module also handles single and multi-byte characters from various codesets including both Extended Unix Code (EUC) and non-EUC codesets.</p> <p>The remainder of this section describes the processing of various STREAMS messages on the read- and write-side.</p>
Read-side Behavior	<p>Various types of STREAMS messages are processed as follows:</p> <p><code>M_BREAK</code> Depending on the state of the <code>BRKINT</code> flag, either an interrupt signal is generated or the message is treated as if it were an <code>M_DATA</code> message containing a single ASCII NUL character when this message is received.</p> <p><code>M_DATA</code> This message is normally processed using the standard <code>termio</code> input processing. If the <code>ICANON</code> flag is set, a single input record (“line”) is accumulated in an internal buffer and sent upstream when a line-terminating character is received. If the <code>ICANON</code> flag is not set, other input processing is performed and the processed data are passed upstream.</p> <p>If output is to be stopped or started as a result of the arrival of characters (usually <code>CNTRL-Q</code> and <code>CNTRL-S</code>), <code>M_STOP</code> and <code>M_START</code> messages are sent downstream. If the <code>IXOFF</code> flag is set and input is to be stopped or started as a result of flow-control considerations, <code>M_STOP_I</code> and <code>M_START_I</code> messages are sent downstream.</p> <p><code>M_DATA</code> messages are sent downstream, as necessary, to perform echoing.</p>

ldterm(7M)

Write-side Behavior

If a signal is to be generated, an `M_FLUSH` message with a flag byte of `FLUSHR` is placed on the read queue. If the signal is also to flush output, an `M_FLUSH` message with a flag byte of `FLUSHW` is sent downstream.

All other messages are passed upstream unchanged.

Various types of STREAMS messages are processed as follows:

- | | |
|-----------------------|--|
| <code>M_FLUSH</code> | The write queue of the module is flushed of all its data messages and the message is passed downstream. |
| <code>M_IOCTL</code> | The function of this <code>ioctl</code> is performed and the message is passed downstream in most cases. The <code>TCFLSH</code> and <code>TCXONC</code> <code>ioctls</code> can be performed entirely in the <code>ldterm</code> module, so the reply is sent upstream and the message is not passed downstream. |
| <code>M_DATA</code> | If the <code>OPOST</code> flag is set, or both the <code>XCASE</code> and <code>ICANON</code> flags are set, output processing is performed and the processed message is passed downstream along with any <code>M_DELAY</code> messages generated. Otherwise, the message is passed downstream without change. |
| <code>M_CTL</code> | If the size of the data buffer associated with the message is the size of <code>struct iocblk</code> , <code>ldterm</code> will perform functional negotiation to determine where the termio(7I) processing is to be done. If the command field of the <code>iocblk</code> structure (<code>ioc_cmd</code>) is set to <code>MC_NO_CANON</code> , the input canonical processing normally performed on <code>M_DATA</code> messages is disabled and those messages are passed upstream unmodified. (This is for the use of modules or drivers that perform their own input processing, such as a pseudo-terminal in <code>TIOCREMOTE</code> mode connected to a program that performs this processing). If the command is <code>MC_DO_CANON</code> , all input processing is enabled. If the command is <code>MC_PART_CANON</code> , then an <code>M_DATA</code> message containing a <code>termios</code> structure is expected to be attached to the original <code>M_CTL</code> message. The <code>ldterm</code> module will examine the <code>iflag</code> , <code>oflag</code> , and <code>lflag</code> fields of the <code>termios</code> structure and from that point on, will process only those flags that have not been turned ON. If none of the above commands are found, the message is ignored. In any case, the message is passed upstream. |
| <code>M_FLUSH</code> | The read queue of the module is flushed of all its data messages and all data in the record being accumulated are also flushed. The message is passed upstream. |
| <code>M_IOCACK</code> | The data contained within the message, which is to be returned to the process, are augmented if necessary, and the message is passed upstream. |

All other messages are passed downstream unchanged.

IOCTLS	<p>The <code>ldterm</code> module processes the following <code>TRANSPARENT</code> ioctls. All others are passed downstream.</p> <p>TCGETS/TCGETA The message is passed downstream. If an acknowledgment is seen, the data provided by the driver and modules downstream are augmented and the acknowledgement is passed upstream.</p> <p>TCSETS/TCSETSW/TCSETSF/TCSETA/TCSETAW/TCSETAF The parameters that control the behavior of the <code>ldterm</code> module are changed. If a mode change requires options at the stream head to be changed, an <code>M_SETOPTS</code> message is sent upstream. If the <code>ICANON</code> flag is turned on or off, the read mode at the stream head is changed to message-nondiscard or byte-stream mode, respectively. If the <code>TOSTOP</code> flag is turned on or off, the <code>tostop</code> mode at the stream head is turned on or off, respectively. In any case, <code>ldterm</code> passes the <code>ioctl</code> on downstream for possible additional processing.</p> <p>TCFLSH If the argument is 0, an <code>M_FLUSH</code> message with a flag byte of <code>FLUSHR</code> is sent downstream and placed on the read queue. If the argument is 1, the write queue is flushed of all its data messages and an <code>M_FLUSH</code> message with a flag byte of <code>FLUSHW</code> is sent upstream and downstream. If the argument is 2, the write queue is flushed of all its data messages and an <code>M_FLUSH</code> message with a flag byte of <code>FLUSHRW</code> is sent downstream and placed on the read queue.</p> <p>TCXONC If the argument is 0 and output is not already stopped, an <code>M_STOP</code> message is sent downstream. If the argument is 1 and output is stopped, an <code>M_START</code> message is sent downstream. If the argument is 2 and input is not already stopped, an <code>M_STOPI</code> message is sent downstream. If the argument is 3 and input is stopped, an <code>M_STARTI</code> message is sent downstream.</p> <p>TCSBRK The message is passed downstream, so the driver has a chance to drain the data and then send an <code>M_IOCACK</code> message upstream.</p> <p>EUC_WSET This call takes a pointer to an <code>euclioc</code> structure, and uses it to set the EUC line discipline's local definition for the code set widths to be used for subsequent operations. Within the stream, the line discipline may optionally notify other modules of this setting using <code>M_CTL</code> messages. When this call is received and the <code>euclioc</code> structure contains valid data, the line discipline changes into EUC handling mode once the <code>eucliocdata</code> is completely transferred to an internal data structure.</p> <p>EUC_WGET This call takes a pointer to an <code>euclioc</code> structure, and returns in it the EUC code set widths currently in use by the EUC line discipline. If the current codeset of the line discipline is not an EUC one, the result is meaningless.</p>
SEE ALSO	<p><code>termios(3C)</code>, <code>console(7D)</code>, <code>termio(7I)</code></p> <p><i>STREAMS Programming Guide</i></p>

llc1(7D)

NAME	llc1 – Logical Link Control Protocol Class 1 Driver
SYNOPSIS	<pre>#include <sys/stropts.h> #include <sys/ethernet.h> #include <sys/dlpi.h> #include <sys/llc1.h></pre>
DESCRIPTION	<p>The llc1 driver is a multi-threaded, loadable, clonable, STREAMS multiplexing driver supporting the connectionless Data Link Provider Interface, dlpi(7P), implementing IEEE 802.2 Logical Link Control Protocol Class 1 over a STREAM to a MAC level driver. Multiple MAC level interfaces installed within the system can be supported by the driver. The llc1 driver provides basic support for the LLC1 protocol. Functions provided include frame transmit and receive, XID, and TEST, multicast support, and error recovery and reporting.</p> <p>The cloning, character-special device, /dev/llc1, is used to access all LLC1 controllers configured under llc1.</p> <p>The llc1 driver is a “Style 2” Data Link Service provider. All messages of types M_PROTO and M_PCPROTO are interpreted as DLPI primitives. An explicit DL_ATTACH_REQ message by the user is required to associate the opened stream with a particular device (ppa). The ppa ID is interpreted as an unsigned long and indicates the corresponding device instance (unit) number. An error (DL_ERROR_ACK) is returned by the driver if the ppa field value does not correspond to a valid device instance number for this system.</p> <p>The values returned by the driver in the DL_INFO_ACK primitive in response to the DL_INFO_REQ from the user are as follows:</p> <ul style="list-style-type: none">■ The maximum Service Data UNIT (SDU) is derived from the MAC layer linked below the driver. In the case of an Ethernet driver, the SDU will be 1497.■ The minimum SDU is 0.■ The MAC type is DL_CSMACD or DL_TPR as determined by the driver linked under llc1. If the driver reports that it is DL_ETHER, it will be changed to DL_CSMACD; otherwise the type is the same as the MAC type.■ The sap length value is -1, meaning the physical address component is followed immediately by a 1-octet sap component within the DLSAP address.■ The service mode is DL_CLDLS.■ The MAC type is DL_CSMACD or DL_TPR as determined by the driver linked under llc1. If the driver reports that it is DL_ETHER, it will be changed to DL_CSMACD; otherwise the type is the same as the MAC type.■ The dlsap address length is 7.■ No optional quality of service (QOS) support is included at present, so the QOS fields should be initialized to 0.■ The DLPI version is DL_VERSION_2.

- The provider style is `DL_STYLE2`.
- The broadcast address value is the broadcast address returned from the lower level driver.

Once in the `DL_ATTACHED` state, the user must send a `DL_BIND_REQ` to associate a particular Service Access Point (SAP) with the stream. The `llc1` driver interprets the `sap` field within the `DL_BIND_REQ` as an IEEE 802.2 “SAP,” therefore valid values for the `sap` field are in the `[0-0xFF]` range with only even values being legal.

The `llc1` driver DLSAP address format consists of the 6-octet physical (e.g., Ethernet) address component followed immediately by the 1-octet `sap` (type) component producing a 7-octet DLSAP address. Applications should *not* hard-code to this particular implementation-specific DLSAP address format, but use information returned in the `DL_INFO_ACK` primitive to compose and decompose DLSAP addresses. The `sap` length, full DLSAP length, and `sap`/physical ordering are included within the `DL_INFO_ACK`. The physical address length can be computed by subtracting the absolute value of the `sap` length from the full DLSAP address length or by issuing the `DL_PHYS_ADDR_REQ` to obtain the current physical address associated with the stream.

Once in the `DL_BOUND` state, the user may transmit frames on the LAN by sending `DL_UNITDATA_REQ` messages to the `llc1` driver. The `llc1` driver will route received frames up all open and bound streams having a `sap` which matches the IEEE 802.2 DSAP as `DL_UNITDATA_IND` messages. Received frames are duplicated and routed up multiple open streams if necessary. The DLSAP address contained within the `DL_UNITDATA_REQ` and `DL_UNITDATA_IND` messages consists of both the `sap` (type) and physical (Ethernet) components.

In addition to the mandatory, connectionless DLPI message set, the driver additionally supports the following primitives:

The `DL_ENABMULTI_REQ` and `DL_DISABMULTI_REQ` primitives enable/disable reception of specific multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. These primitives are accepted by the driver in any driver state that is valid while still being attached to the `ppa`.

The `DL_PHYS_ADDR_REQ` primitive returns the 6-octet physical address currently associated (attached) to the stream in the `DL_PHYS_ADDR_ACK` primitive. This primitive is valid only in states following a successful `DL_ATTACH_REQ`.

The `DL_SET_PHYS_ADDR_REQ` primitive changes the 6-octet physical address currently associated (attached) to this stream. Once changed, all streams subsequently opened and attached to this device will obtain this new physical address. Once changed, the physical address will remain set until this primitive is used to change the physical address again or the system is rebooted, whichever occurs first.

The `DL_XID_REQ`/`DL_TEST_REQ` primitives provide the means for a user to issue an LLC XID or TEST request message. A response to one of these messages will be in the form of a `DL_XID_CON`/`DL_TEST_CON` message.

llc1(7D)

The `DL_XID_RES`/`DL_TEST_RES` primitives provide a way for the user to respond to the receipt of an `XID` or `TEST` message that was received as a `DL_XID_IND`/`DL_TEST_IND` message.

`XID` and `TEST` will be automatically processed by `llc1` if the `DL_AUTO_XID`/`DL_AUTO_TEST` bits are set in the `DL_BIND_REQ`.

FILES /dev/llc1 cloning, character-special device

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`, `dlpi(7P)`

NAME	llc2 – Class II logical link control driver								
DESCRIPTION	<p>The llc2 logical link control driver interfaces network software (NetBIOS, SNA, OSI, and so on) running under the Solaris operating environment to a physical LAN network controlled by one of the supported communications adapters. The llc2 driver, which appears as a STREAMS driver to the network software, resides in the kernel and is accessed by standard UNIX STREAMS functions.</p> <p>This version of the llc2 driver includes support for both connectionless and connection-oriented logical link control class II (llc2) operations for Ethernet, Token Ring, and FDDI adapters when accessed through the appropriate Solaris MAC layer driver. The Data Link Provider Interface (DLPI) to the llc2 driver enables multiple and different protocol stacks, (including NetBIOS and SNA), to operate simultaneously over one or more local area networks.</p> <p>To start the llc2 driver by default, rename file /etc/llc2/llc2_start.default to /etc/llc2/llc2_start. This allows the /etc/rc2.d/S40llc2 script to build up the configuration file for each ppa interface in /etc/llc2/default/llc2.* and start llc2 on each interface. To verify the configuration files, manually run /usr/lib/llc2/llc2_autoconfig.</p> <p>For more information on the llc2 driver, see the IEEE standard <i>802.2 Logical Link Control</i>.</p>								
OBTAINING LLC2 STATISTICS	<p>You can obtain LLC2 statistics or reset the statistics counter to zero using the ILD_LLCC2 ioctl. The ILD_LLCC2 ioctl has a number of subcommands. The following retrieve LLC2 statistics:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">Name</th> <th style="text-align: center;">Function</th> </tr> </thead> <tbody> <tr> <td>LLC2_GET_STA_STATS</td> <td>Get station statistics</td> </tr> <tr> <td>LLC2_GET_SAP_STATS</td> <td>Get SAP statistics</td> </tr> <tr> <td>LLC2_GET_CON_STATS</td> <td>Get connection statistics</td> </tr> </tbody> </table>	Name	Function	LLC2_GET_STA_STATS	Get station statistics	LLC2_GET_SAP_STATS	Get SAP statistics	LLC2_GET_CON_STATS	Get connection statistics
Name	Function								
LLC2_GET_STA_STATS	Get station statistics								
LLC2_GET_SAP_STATS	Get SAP statistics								
LLC2_GET_CON_STATS	Get connection statistics								
LLC2_GET_STA_STATS	<p>The structure used depends on the subcommand sent.</p> <p>The LLC2_GET_STA_STATS command retrieves statistics on a particular Physical Point of Attachment (PPA).</p> <p>When sending the LLC2_GET_STA_STATS command, the <i>llc2GetStaStats</i> structure is used:</p> <pre>typedef struct llc2GetStaStats { uint_t ppa; uint_t cmd; uchar_t clearFlag; uchar_t state; ushort_t numSaps; uchar_t saps[LLC2_MAX_SAPS]; }</pre>								

llc2(7D)

```
uint_t nullSapXidCmdRcvd;
uint_t nullSapXidRspSent;
uint_t nullSapTestCmdRcvd;
uint_t nullSapTestRspSent;
uint_t outOfState;
uint_t allocFail;
uint_t protocolError;
} llc2GetStaStats_t;
```

The members of the structure are:

Member	Description
cmd	LLC2_GET_STA_STATS
clearFlag	Clear counters flag. Set this to 0 to retrieve statistics and to 1 to reset all counters to 0.
state	Station component state. Possible values are ?????
numSaps	Number of active SAPs in the saps array
saps	An array of active SAP values
nullSapXidCmdRcvd	Number of NULL SAP XID commands received
nullSapXidRspSent	Number of NULL SAP XID responses sent
nullSapTestCmdRcvd	Number of NULL SAP TEST commands received
nullSapTestRspSent	Number of NULL SAP TEST responses sent
outOfState	Number of invalid events received
allocFail	Number of buffer allocation failures
protocolError	Number of protocol errors

LLC2_GET_SAP_STATS

The LLC2_GET_SAP_STATS command retrieves statistics related to a particular SAP. When sending the LLC2_GET_SAP_STATS command, the *llc2GetSapStats* structure is used:

```
typedef struct llc2GetSapStats {
uint_t ppa;
uint_t cmd;
uchar_t sap;
uchar_t clearFlag;
uchar_t state;
uint_t numCons;
ushort_t cons[LLC2_MAX_CONS];
uint_t xidCmdSent;
uint_t xidCmdRcvd;
uint_t xidRspSent;
uint_t xidRspRcvd;
uint_t testCmdSent;
uint_t testCmdRcvd;
uint_t testRspSent;
```

```

uint_t testRspRcvd;
uint_t uiSent;
uint_t uiRcvd;
uint_t outOfState;
uint_t allocFail;
uint_t protocolError;
} llc2GetSapStats_t;

```

The members are:

Member	Description
ppa	Physical Point of Attachment number
cmd	LLC2_GET_SAP_STATS
sap	SAP value
clearFlag	Clear counters flag. Set this to 0 to retrieve statistics and to 1 to reset all counters to 0.
state	SAP component state
numCons	Number of active connections in the cons array
cons	Array of active connection indexes
xidCmdSent	Number of XID commands sent
xidCmdRcvd	Number of XID responses received
xidRspSent	Number of XID responses sent
xidRspRcvd	Number of XID responses received
testCmdSent	Number of TEST commands sent
testCmdRcvd	Number of TEST commands received
testRspSent	Number of TEST responses sent
testRspRcvd	Number of TEST responses received
uiSent	Number of UI frames sent
uiRcvd	Number of UI frames received
outOfState	Number of invalid events received
allocFail	Number of buffer allocation failures
protocolError	Number of protocol errors

LLC2_GET_CON_STATS

The LLC2_GET_CON_STATS command retrieves statistics related to a particular connection component. When sending the LLC2_GET_CON_STATS command, the *llc2GetConStats* structure is used:

llc2(7D)

```
typedef struct llc2GetConStats {
    uint_t ppa;
    uint_t cmd;
    uchar_t sap;
    ushort_t con;
    uchar_t clearFlag;
    uchar_t stateOldest;
    uchar_t stateOlder;
    uchar_t stateOld;
    uchar_t state;
    ushort_t sid;
    dlsap_t rem;
    ushort_t flag;
    uchar_t dataFlag;
    uchar_t k;
    uchar_t vs;
    uchar_t vr;
    uchar_t nrRcvd;
    ushort_t retryCount;
    uint_t numToBeAcked;
    uint_t numToResend;
    uint_t macOutSave;
    uint_t macOutDump;
    uchar_t timerOn;
    uint_t iSent;
    uint_t iRcvd;
    uint_t frmrSent;
    uint_t frmrRcvd;
    uint_t rrSent;
    uint_t rrRcvd;
    uint_t rnrSent;
    uint_t rnrRcvd;
    uint_t rejSent;
    uint_t rejRcvd;
    uint_t sabmeSent;
    uint_t sabmeRcvd;
    uint_t uaSent;
    uint_t uaRcvd;
    uint_t discSent;
    uint_t outOfState;
    uint_t allocFail;
    uint_t protocolError;
    uint_t localBusy;
    uint_t remoteBusy;
    uint_t maxRetryFail;
    uint_t ackTimerExp;
    uint_t pollTimerExp;
    uint_t rejTimerExp;
    uint_t remBusyTimerExp;
    uint_t inactTimerExp;
    uint_t sendAckTimerExp;
} llc2GetConStats_t;
```

The members of the structure are:

Member	Description
ppa	Physical Point of Attachment number
cmd	LLC2_GET_CON_STATS
sap	SAP value
con	Connection index
clearFlag	Clear counters flag. Set this to 0 to retrieve statistics and to 1 to reset all counters to 0.
stateOldest, stateOlder, stateOld, state	The four previous dlpi states of the connection
sid	SAP value and connection index
dlsap_t rem	Structure containing the remote MAC address and SAP
flag	Connection component processing flag
dataFlag	DATA_FLAG
k	transmit window size
vs	Sequence number of the next I-frame to send
vr	Sequence number of the next I-frame expected
nrRcvd	Sequence number of the last I-frame acknowledged by the remote node
retryCount	Number of timer expirations
numToBeAcked	Number of outbound I-frames to be acknowledged
numToResend	Number of outbound I-frames to be re-sent
macOutSave	Number of outbound I-frames held by the MAC driver to be saved on return to LLC2
macOutDump	Number of outbound I-frames held by the MAC driver to be dumped on return to LLC2
timerOn	Timer activity flag
iSent	Number of I-frames sent
iRcvd	Number of I-frames received
frmrSent	Number of frame rejects sent
frmrRcvd	Number of frame rejects received
rrSent	Number of RRs sent
rrRcvd	Number of RRs received

llc2(7D)

Member	Description
rnrRcvd	Number of RNRs received
rejSent	Number of rejects sent
rejRcvd	Number of rejects received
sabmeSent	Number of SABMEs sent
sabmeRcvd	Number of SABMEs received
uaSent	Number of UAs sent
uaRcvd	Number of UAs received
discSent	Number of DISCs sent
outOfState	Number of invalid events received
allocFail	Number of buffer allocation failures
protocolError	Number of protocol errors
localBusy	Number of times in a local busy state
remoteBusy	Number of times in a remote busy state
maxRetryFail	Number of failures due to reaching maxRetry
ackTimerExp	Number of ack timer expirations
pollTimerExp	Number of P-timer expirations
rejTimerExp	Number of reject timer expirations
remBusyTimerExp	Number of remote busy timer expirations
inactTimerExp	Number of inactivity timer expirations
sendAckTimerExp	Number of send ack timer expirations

FILES /dev/llc2 Clone device used to access the driver
 /etc/llc2/default/llc2.? configuration files
 (One file per ppa interface.)

ATTRIBUTES See attributes(5) for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWllc

SEE ALSO llc2_autoconfig(1), llc2_config(1), llc2(4)

NAME	lockstat – DTrace kernel lock instrumentation provider
DESCRIPTION	<p>The <code>lockstat</code> driver is a DTrace dynamic tracing provider that performs dynamic instrumentation for locking primitives in the Solaris kernel.</p> <p>The <code>lockstat</code> provider makes probes available that you can use to discern lock contention statistics, or to understand virtually any aspect of locking behavior inside the operating system kernel. The <code>lockstat(1M)</code> command is implemented as a DTrace consumer that uses the <code>lockstat</code> provider to gather raw data.</p> <p>The <code>lockstat</code> driver is not a public interface and you access the instrumentation offered by this provider through DTrace. Refer to the <i>Solaris Dynamic Tracing Guide</i> for a description of the public documented interfaces available for the DTrace facility and the probes offered by the <code>lockstat</code> provider.</p>
ATTRIBUTES	See <code>attributes(5)</code> for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdtrp
Interface Stability	Private

SEE ALSO `dtrace(1M)`, `lockstat(1M)`, `attributes(5)`, `dtrace(7D)`

Solaris Dynamic Tracing Guide

lofi(7D)

NAME	lofi – Loopback file driver				
DESCRIPTION	<p>The <code>lofi</code> file driver exports a file as a block device. Reads and writes to the block device are translated to reads and writes on the underlying file. This is useful when the file contains a file system image. Exporting it as a block device through the <code>lofi</code> file driver allows normal system utilities to operate on the image through the block device (like <code>fstyp(1M)</code>, <code>fsck(1M)</code>, and <code>mount(1M)</code>). This is useful for accessing CD-ROM and FAT floppy images. See <code>lofiadm(1M)</code> for examples.</p> <p>File block device entries are contained in <code>/dev/lofi</code>, while <code>/dev/rlofi</code> contains the character (or raw) device entries. Entries are in the form of decimal numbers which are assigned through <code>lofiadm(1M)</code>. When created, these device entries are owned by <code>root</code>, in group <code>sys</code>, and have permissions <code>0600</code>. While ownership, group, and permission settings can be altered, there are possible ramifications. See <code>lofiadm(1M)</code> for more information.</p>				
FILES	<p><code>/dev/lofictl</code> Master control device</p> <p><code>/dev/lofi/n</code> Block device for file <i>n</i></p> <p><code>/dev/rlofi/n</code> Character device for file <i>n</i></p> <p><code>/kernel/drv/lofi</code> 32-bit driver</p> <p><code>/kernel/drv/lofi.conf</code> Driver configuration file. (Should not be altered.)</p> <p><code>/kernel/drv/sparcv9/lofi</code> 64-bit driver</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsr, SUNWcarx.u</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsr, SUNWcarx.u
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	SUNWcsr, SUNWcarx.u				
SEE ALSO	<code>lofiadm(1M)</code> , <code>fsck(1M)</code> , <code>fstyp(1M)</code> , <code>mount(1M)</code> , <code>newfs(1M)</code> , <code>attributes(5)</code> , <code>lofs(7FS)</code>				
NOTES	<p>Just as you would not directly access a disk device that has mounted file systems, you should not access a file associated with a block device except through the <code>lofi</code> file driver.</p> <p>For compatibility purposes, a raw device is also exported along with the block device. For example, <code>newfs(1M)</code> requires one.</p>				

NAME	lofs – loopback virtual file system
SYNOPSIS	<pre>#include <sys/param.h> #include <sys/mount.h> int mount (const char* <i>dir</i>, const char* <i>virtual</i>, int <i>mflag</i>, <i>lofs</i>, NULL, 0);</pre>
DESCRIPTION	<p>The loopback file system device allows new, virtual file systems to be created, which provide access to existing files using alternate pathnames. Once the virtual file system is created, other file systems can be mounted within it, without affecting the original file system. However, file systems which are subsequently mounted onto the original file system <i>are</i> visible to the virtual file system, unless or until the corresponding mount point in the virtual file system is covered by a file system mounted there.</p> <p><i>virtual</i> is the mount point for the virtual file system. <i>dir</i> is the pathname of the existing file system. <i>mflag</i> specifies the mount options; the <code>MS_DATA</code> bit in <i>mflag</i> must be set. If the <code>MS_RDONLY</code> bit in <i>mflag</i> is not set, accesses to the loop back file system are the same as for the underlying file system. Otherwise, all accesses in the loopback file system will be read-only. All other <code>mount(2)</code> options are inherited from the underlying file systems.</p> <p>A loopback mount of <code>''</code> onto <code>/tmp/newroot</code> allows the entire file system hierarchy to appear as if it were duplicated under <code>/tmp/newroot</code>, including any file systems mounted from remote NFS servers. All files would then be accessible either from a pathname relative to <code>''</code> or from a pathname relative to <code>/tmp/newroot</code> until such time as a file system is mounted in <code>/tmp/newroot</code>, or any of its subdirectories.</p> <p>Loopback mounts of <code>''</code> can be performed in conjunction with the <code>chroot(2)</code> system call, to provide a complete virtual file system to a process or family of processes.</p> <p>Recursive traversal of loopback mount points is not allowed. After the loopback mount of <code>/tmp/newroot</code>, the file <code>/tmp/newroot/tmp/newroot</code> does not contain yet another file system hierarchy; rather, it appears just as <code>/tmp/newroot</code> did before the loopback mount was performed (for example, as an empty directory).</p>
Examples	<p>lofs file systems are mounted using:</p> <pre>mount -F lofs /tmp /mnt</pre>
SEE ALSO	lofiadm(1M), mount(1M), chroot(2), mount(2), sysfs(2), vfstab(4), lofi(7D)
WARNINGS	Loopback mounts must be used with care; the potential for confusing users and applications is enormous. A loopback mount entry in <code>/etc/vfstab</code> must be placed after the mount points of both directories it depends on. This is most easily accomplished by making the loopback mount entry the last in <code>/etc/vfstab</code> .

log(7D)

NAME	log – interface to STREAMS error logging and event tracing
SYNOPSIS	<pre>#include <sys/strlog.h> #include <sys/log.h></pre>
DESCRIPTION	<p>log is a STREAMS software device driver that provides an interface for console logging and for the STREAMS error logging and event tracing processes (see <code>strerr(1M)</code>, and <code>strace(1M)</code>). log presents two separate interfaces: a function call interface in the kernel through which STREAMS drivers and modules submit log messages; and a set of <code>ioctl(2)</code> requests and STREAMS messages for interaction with a user level console logger, an error logger, a trace logger, or processes that need to submit their own log messages.</p>
Kernel Interface	<p>log messages are generated within the kernel by calls to the function <code>strlog()</code>:</p> <pre>strlog(short mid, short sid, char level, ushort_t flags, char *fmt, unsigned arg1 . . .);</pre> <p>Required definitions are contained in <code><sys/strlog.h></code>, <code><sys/log.h></code>, and <code><sys/syslog.h></code>. <i>mid</i> is the STREAMS module id number for the module or driver submitting the log message. <i>sid</i> is an internal sub-id number usually used to identify a particular minor device of a driver. <i>level</i> is a tracing level that allows for selective screening out of low priority messages from the tracer. <i>flags</i> are any combination of <code>SL_ERROR</code> (the message is for the error logger), <code>SL_TRACE</code> (the message is for the tracer), <code>SL_CONSOLE</code> (the message is for the console logger), <code>SL_FATAL</code> (advisory notification of a fatal error), and <code>SL_NOTIFY</code> (request that a copy of the message be mailed to the system administrator). <i>fmt</i> is a <code>printf(3C)</code> style format string, except that <code>%s</code>, <code>%e</code>, <code>%E</code>, <code>%g</code>, and <code>%G</code> conversion specifications are not handled. Up to <code>NLOGARGS</code> (in this release, three) numeric or character arguments can be provided.</p>
User Interface	<p>log is implemented as a cloneable device, it clones itself without intervention from the system clone device. Each open of <code>/dev/log</code> obtains a separate stream to log. In order to receive log messages, a process must first notify log whether it is an error logger, trace logger, or console logger using a STREAMS <code>I_STR</code> <code>ioctl</code> call (see below). For the console logger, the <code>I_STR</code> <code>ioctl</code> has an <code>ic_cmd</code> field of <code>I_CONSLOG</code>, with no accompanying data. For the error logger, the <code>I_STR</code> <code>ioctl</code> has an <code>ic_cmd</code> field of <code>I_ERRLOG</code>, with no accompanying data. For the trace logger, the <code>ioctl</code> has an <code>ic_cmd</code> field of <code>I_TRCLOG</code>, and must be accompanied by a data buffer containing an array of one or more <code>struct trace_ids</code> elements.</p> <pre>struct trace_ids { short ti_mid; short ti_sid; char ti_level; };</pre>

Each `trace_ids` structure specifies a *mid*, *sid*, and *level* from which messages will be accepted. `strlog(9F)` will accept messages whose *mid* and *sid* exactly match those in the `trace_ids` structure, and whose level is less than or equal to the level given in the `trace_ids` structure. A value of `-1` in any of the fields of the `trace_ids` structure indicates that any value is accepted for that field.

Once the logger process has identified itself using the `ioctl` call, `log` will begin sending up messages subject to the restrictions noted above. These messages are obtained using the `getmsg(2)` function. The control part of this message contains a `log_ctl` structure, which specifies the *mid*, *sid*, *level*, *flags*, time in ticks since boot that the message was submitted, the corresponding time in seconds since Jan. 1, 1970, a sequence number, and a priority. The time in seconds since 1970 is provided so that the date and time of the message can be easily computed, and the time in ticks since boot is provided so that the relative timing of `log` messages can be determined.

```
struct log_ctl {
    short mid;
    short sid;
    char level;      /* level of message for tracing */
    short flags;    /* message disposition */
#if defined(_LP64) || defined(_I32LPx)
    clock32_t ltime; /* time in machine ticks since boot */
    time32_t ttime; /* time in seconds since 1970 */
#else
    clock_t ltime;
    time_t ttime;
#endif
    int seq_no;     /* sequence number */
    int pri;        /* priority = (facility|level) */
};
```

The priority consists of a priority code and a facility code, found in `<sys/syslog.h>`. If `SL_CONSOLE` is set in *flags*, the priority code is set as follows: If `SL_WARN` is set, the priority code is set to `LOG_WARNING`; If `SL_FATAL` is set, the priority code is set to `LOG_CRIT`; If `SL_ERROR` is set, the priority code is set to `LOG_ERR`; If `SL_NOTE` is set, the priority code is set to `LOG_NOTICE`; If `SL_TRACE` is set, the priority code is set to `LOG_DEBUG`; If only `SL_CONSOLE` is set, the priority code is set to `LOG_INFO`. Messages originating from the kernel have the facility code set to `LOG_KERN`. Most messages originating from user processes will have the facility code set to `LOG_USER`.

Different sequence numbers are maintained for the error and trace logging streams, and are provided so that gaps in the sequence of messages can be determined (during times of high message traffic some messages may not be delivered by the logger to avoid hogging system resources). The data part of the message contains the unexpanded text of the format string (null terminated), followed by `NLOGARGS` words for the arguments to the format string, aligned on the first word boundary following the format string.

A process may also send a message of the same structure to `log`, even if it is not an error or trace logger. The only fields of the `log_ctl` structure in the control part of the message that are accepted are the *level*, *flags*, and *pri* fields; all other fields are filled in

log(7D)

by log before being forwarded to the appropriate logger. The data portion must contain a null terminated format string, and any arguments (up to NLOGARGS) must be packed, 32-bits each, on the next 32-bit boundary following the end of the format string.

ENXIO is returned for I_TRCLOG ioctls without any trace_ids structures, or for any unrecognized ioctl calls. The driver silently ignores incorrectly formatted log messages sent to the driver by a user process (no error results).

Processes that wish to write a message to the console logger may direct their output to /dev/conslog, using either write(2) or putmsg(2).

Driver Configuration

The following driver configuration properties may be defined in the log.conf file.

msgid=1 If msgid=1, each message will be preceded by a message ID as described in syslogd(1M).

msgid=0 If msgid=0, message IDs will not be generated. This property is unstable and may be removed in a future release.

EXAMPLES

EXAMPLE 1 I_ERRLOG registration.

```
struct strioctl ioc;
ioc.ic_cmd = I_ERRLOG;
ioc.ic_timeout = 0;            /* default timeout (15 secs.) */
ioc.ic_len = 0;
ioc.ic_dp = NULL;
ioctl(log, I_STR, &ioc);
```

EXAMPLE 2 I_TRCLOG registration.

```
struct trace_ids tid[2];
tid[0].ti_mid = 2;
tid[0].ti_sid = 0;
tid[0].ti_level = 1;
tid[1].ti_mid = 1002;
tid[1].ti_sid = -1;            /* any sub-id will be allowed */
tid[1].ti_level = -1;        /* any level will be allowed */
ioc.ic_cmd = I_TRCLOG;
ioc.ic_timeout = 0;
ioc.ic_len = 2 * sizeof(struct trace_ids);
ioc.ic_dp = (char *)tid;
ioctl(log, I_STR, &ioc);
```

Example of submitting a log message (no arguments):

```
struct strbuf ctl, dat;
struct log_ctl lc;
char *message = "Don't forget to pick up some milk
                  on the way home";
ctl.len = ctl.maxlen = sizeof(lc);
ctl.buf = (char *)&lc;
dat.len = dat.maxlen = strlen(message);
dat.buf = message;
```


EXAMPLE 2 I_TRCLOG registration. (Continued)

```
lc.level = 0;
lc.flags = SL_ERROR|SL_NOTIFY;
putmsg(log, &ctl, &dat, 0);
```

FILES

/dev/log	Log driver.
/dev/conslog	Write only instance of the log driver, for console logging.
/kernel/drv/log.conf	Log configuration file.

SEE ALSO strace(1M), strerr(1M), intro(3), getmsg(2), ioctl(2), putmsg(2), write(2), printf(3C), strlog(9F)

STREAMS Programming Guide

logi(7D)

NAME logi – LOGITECH Bus Mouse device interface

SYNOPSIS /dev/logi

DESCRIPTION The logi driver supports the LOGITECH Bus Mouse. It allows applications to obtain information about the mouse's movements and the status of its buttons. The data is read in the Five Byte Packed Binary Format, also called MSC format.

FILES /dev/logi

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`

NAME	lp – driver for parallel port				
SYNOPSIS	<pre>include <sys/bpp_io.h> fd = open("/dev/lp<i>n</i>", flags);</pre>				
DESCRIPTION	The lp driver provides the interface to the parallel ports used by printers for x86 based systems. The lp driver is implemented as a STREAMS device.				
IOCTLS	<table border="0"> <tr> <td style="vertical-align: top;">BPPIOC_TESTIO</td> <td>Test transfer readiness. This command checks to see if a read or write transfer would succeed based on pin status. If a transfer would succeed, 0 is returned. If a transfer would fail, -1 is returned, and <code>errno</code> is set to <code>EIO</code>. The error status can be retrieved using the <code>BPPIOC_GETERR</code> <code>ioctl()</code> call.</td> </tr> <tr> <td style="vertical-align: top;">BPPIOC_GETERR</td> <td>Get last error status. The argument is a pointer to a <code>struct bpp_error_status</code>. See below for a description of the elements of this structure. This structure indicates the status of all the appropriate status bits at the time of the most recent error condition during a <code>read(2)</code> or <code>write(2)</code> call, or the status of the bits at the most recent <code>BPPIOC_TESTIO</code> <code>ioctl(2)</code> call. The application can check transfer readiness without attempting another transfer using the <code>BPPIOC_TESTIO</code> <code>ioctl()</code>.</td> </tr> </table>	BPPIOC_TESTIO	Test transfer readiness. This command checks to see if a read or write transfer would succeed based on pin status. If a transfer would succeed, 0 is returned. If a transfer would fail, -1 is returned, and <code>errno</code> is set to <code>EIO</code> . The error status can be retrieved using the <code>BPPIOC_GETERR</code> <code>ioctl()</code> call.	BPPIOC_GETERR	Get last error status. The argument is a pointer to a <code>struct bpp_error_status</code> . See below for a description of the elements of this structure. This structure indicates the status of all the appropriate status bits at the time of the most recent error condition during a <code>read(2)</code> or <code>write(2)</code> call, or the status of the bits at the most recent <code>BPPIOC_TESTIO</code> <code>ioctl(2)</code> call. The application can check transfer readiness without attempting another transfer using the <code>BPPIOC_TESTIO</code> <code>ioctl()</code> .
BPPIOC_TESTIO	Test transfer readiness. This command checks to see if a read or write transfer would succeed based on pin status. If a transfer would succeed, 0 is returned. If a transfer would fail, -1 is returned, and <code>errno</code> is set to <code>EIO</code> . The error status can be retrieved using the <code>BPPIOC_GETERR</code> <code>ioctl()</code> call.				
BPPIOC_GETERR	Get last error status. The argument is a pointer to a <code>struct bpp_error_status</code> . See below for a description of the elements of this structure. This structure indicates the status of all the appropriate status bits at the time of the most recent error condition during a <code>read(2)</code> or <code>write(2)</code> call, or the status of the bits at the most recent <code>BPPIOC_TESTIO</code> <code>ioctl(2)</code> call. The application can check transfer readiness without attempting another transfer using the <code>BPPIOC_TESTIO</code> <code>ioctl()</code> .				
Error Pins Structure	<p>This structure and symbols are defined in the include file <code><sys/bpp_io.h></code>:</p> <pre>struct bpp_error_status { char timeout_occurred; /* Not use */ char bus_error; /* Not use */ uchar_t pin_status; /* Status of pins which could cause an error */ }; /* Values for pin_status field */ #define BPP_ERR_ERR 0x01 /* Error pin active */ #define BPP_SLCT_ERR 0x02 /* Select pin active */ #define BPP_PE_ERR 0x04 /* Paper empty pin active */</pre> <p>Note: Other pin statuses are defined in <code><sys/bpp_io.h></code>, but <code>BPP_ERR_ERR</code>, <code>BPP_SLCT_ERR</code> and <code>BPP_PE_ERR</code> are the only ones valid for the x86 lp driver.</p>				
ERRORS	<table border="0"> <tr> <td style="vertical-align: top;">EIO</td> <td>A <code>BPPIOC_TESTIO</code> <code>ioctl()</code> call is attempted while a condition exists that would prevent a transfer (such as a peripheral error).</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>An <code>ioctl()</code> is attempted with an invalid value in the command argument.</td> </tr> </table>	EIO	A <code>BPPIOC_TESTIO</code> <code>ioctl()</code> call is attempted while a condition exists that would prevent a transfer (such as a peripheral error).	EINVAL	An <code>ioctl()</code> is attempted with an invalid value in the command argument.
EIO	A <code>BPPIOC_TESTIO</code> <code>ioctl()</code> call is attempted while a condition exists that would prevent a transfer (such as a peripheral error).				
EINVAL	An <code>ioctl()</code> is attempted with an invalid value in the command argument.				
FILES	<code>/platform/i86pc/kernel/drv/lp.conf</code> configuration file for lp driver				

lp(7D)

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `sysbus(4)`, `attributes(5)`, `streamio(7I)`

NOTES A read operation on a bi-directional parallel port is not supported.

NAME m64 – PGX, PGX24, and PGX64 frame buffers device driver

DESCRIPTION The m64 driver is the Sun PGX graphics accelerator device driver.

FILES /dev/fbs/m64* In Device special file

SEE ALSO m64config(1M)

md(7D)

NAME	md – user configurable pseudo device driver
DESCRIPTION	<p>md is a user configurable pseudo device driver that provides disk concatenation, striping, mirroring, RAID5 metadevices, trans metadevices, and hot spare utilities. Trans devices are no longer supported and have been replaced by UFS logging. See <code>mount_ufs(1M)</code>.</p> <p>The block devices access the disk using the system's normal buffering mechanism and are read and written without regard to physical disk records. There is also a "raw" device which provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in one I/O operation; raw I/O is therefore considerably more efficient when many bytes are transmitted. The names of the block devices are found in <code>/dev/md/dsk</code>; the names of the raw devices are found in <code>/dev/md/rdisk</code>. Metadevices have the appearance of whole disks; there are no slices (partitions).</p> <p>I/O requests (such as <code>lseek(2)</code>) to the metadevices must have an offset that is a multiple of 512 bytes (<code>DEV_BSIZE</code>), or the driver returns an <code>EINVAL</code> error. If the transfer length is not a multiple of 512 bytes, the transfer count is rounded up by the driver.</p> <p>The md pseudo device drivers support all disk devices on all Solaris 2.4 or later Solaris systems.</p>
IOCTLS	<p>This section provides a list of the ioctls supported by the metadisk driver.</p> <p>The following ioctls are valid when issued to the raw metadvice, such as <code>/dev/md/rdisk/d0</code>. See <code>dkio(7i)</code> for additional information.</p> <p>DKIOCGGEOM This ioctl is used to get the disk geometry. The metadisk driver fills in the <code>dkg_nhead</code>, <code>dkg_nsect</code>, <code>dkg_rpm</code>, <code>dkg_write_reinstruct</code> and <code>dkg_read_reinstruct</code> from the first component of the metadvice (at <code>metainit</code> time). <code>dkg_ncyl</code> is calculated using the size of the metadvice (reported by <code>metastat</code>) divided by $(dkg_nhead * dkg_nsect)$. The total size is always a multiple of $(dkg_nhead * dkg_nsect)$. If the first component of a metadvice <i>does not</i> start on cylinder number 0, then the <code>dkg_ncyl</code> is increased by one cylinder; because <code>DKIOCGVTOC</code> reports the metadvice as starting on cylinder 1. The side effect here is that it looks like cylinder 0 is not being used, but all the arithmetic works out correctly. If the metadvice is not set up, then <code>ENXIO</code> is returned.</p> <p>DKIOCIINFO When issued to the administrative device or metadvice, this ioctl sets <code>dki_unit</code> to the unit number of the metadvice, <code>dki_ctype</code> to a value of <code>DKC_MD</code>, and <code>dki_partition</code> to 0, because there are no slices.</p>

DKIOCGVTOC This ioctl returns the current vtoc. If one has not been written, then a default vtoc is returned. `v_nparts` is always 1. `v_part[0].p_start` is 0 if the first component of the metadvice starts on cylinder 0. Otherwise, the `p_start` field is the starting sector of cylinder 1. `v_part[0].p_size` is the same as the total size reported by `metastat`.

DKIOCSVTOC This ioctl stores the vtoc in the metadvice state database so it is persistent across reboots.

DIAGNOSTICS

Notice Log Messages

The informative log messages include:

```
md: dnum: Hotspared device dev with dev
```

The first device name listed has been hot spare replaced with the second device name listed.

```
md: dnum: Hotspared device dev(num,num) with dev(num,num)
```

The first device number listed has been hot spare replaced with the second device number listed.

```
md: Could not load misc /dev
```

The named `misc` module is not loadable. It is possibly missing, or something else has been copied over it.

```
md: dnum: no mem for property dev
```

Memory could not be allocated in the `prop_op` entry point.

```
md: db: Parsing error on 'dev'
```

Set command in `/kernel/drv/md.conf` for the `mddb.bootlist <number>` is not in the correct format. `metadb -p` can be run to put the correct set commands into the `/kernel/drv/md.conf` file.

```
md: dnum: dev(num,num) needs maintenance
```

```
md: dnum: dev needs maintenance
```

An I/O or open error has occurred on a device within a mirror causing a component in the mirror to change to the Maintenance state.

```
md: dnum: dev(num,num) last erred md: dnum: dev last erred
```

An I/O or open error has occurred on a device within a mirror and the data is not replicated elsewhere in the mirror. This is causing the component in the mirror to change to the Last Erred state.

Warning Log Messages

The warning log messages include:

md(7D)

```
md: State database is stale
```

This error message comes when there are not enough usable replicas for the state database to be able to update records in the database. All accesses to the metadvice driver will fail. To fix this problem, more replicas need to be added or inaccessible replicas need to be deleted.

```
md: dnum: read error on devmd: dnum: write error on dev
```

A read or write error has occurred on the specified submirror, at the specified device name. This happens if any read or write errors occur on a submirror.

```
md: dnum: read error on dev(num,num)md: dnum: write error on dev(num,num)
```

A read or write error has occurred on the specified submirror, at the specified device number. This happens if any read or write errors occur on a submirror.

```
md: State database commit failed
```

```
md: State database delete failed
```

These messages occur when there have been device errors on components where the state database replicas reside. These errors only occur when more than half of the replicas have had device errors returned to them. For instance, if you have three components with state database replicas and two of the components report errors, then these errors may occur. The state database commit or delete is retried periodically. If a replica is added, then the commit or delete will finish and the system will be operational. Otherwise the system will timeout and panic.

```
md: dnum: Cannot load dev driver
```

Underlying named driver module is not loadable (for example, *sd*, *id*, *xy*, or a third-party driver). This could indicate that the driver module has been removed.

```
md: Open error of hotspare devmd: Open error of hotspare dev(num,num)
```

Named hotspare is not openable, or underlying driver is not loadable.

Panic Log Messages

The panic log messages include:

```
md: dnum: Unknown close typemd: dnum: Unknown open type
```

Metadvice is being opened/closed with an unknown open type (OTYP).

```
md: State database problem
```

Failed metadvice state database commit or delete has been retried the default 100 times.

FILES	<code>/dev/md/dsk/dn</code>	block device (where <i>n</i> is the device number)
	<code>/dev/md/rdisk/dn</code>	raw device (where <i>n</i> is the device number)
	<code>/dev/md/setname/dsk/dn</code>	block device (where <i>setname</i> is the name of the diskset and <i>n</i> is the device number)

/dev/md/ <i>setname</i> / <i>rdsk</i> / <i>dn</i>	raw device (where <i>setname</i> is the name of the diskset and <i>n</i> is the device number)
/dev/md/admin	administrative device
/kernel/drv/md	driver module
/kernel/drv/md.conf	driver configuration file
/kernel/misc/md_stripe	stripe driver misc module
/kernel/misc/md_mirror	mirror driver misc module
/kernel/misc/md_hotspares	hotspares driver misc module
/kernel/misc/md_trans	metatrans driver for UFS logging
/kernel/misc/md_raid	RAID5 driver misc module

SEE ALSO mdmonitord(1M), metaclear(1M), metadb(1M), metadetach(1M), metahs(1M), metainit(1M), metaoffline(1M), metaonline(1M), metaparam(1M), metarecover(1M), metarename(1M), metareplace(1M), metaroot(1M), metassist(1M), metaset(1M), metastat(1M), metasync(1M), metattach(1M), md.cf(4), md.tab(4), attributes(5),

Solaris Volume Manager Administration Guide

NOTES Trans metadevices have been replaced by UFS logging. Existing trans devices are *not* logging--they pass data directly through to the underlying device. See mount_ufs(1M) for more information about UFS logging.

mediator(7D)

NAME	mediator – support for HA configurations consisting of two strings of drives												
DESCRIPTION	<p>Beginning with a prior version, Solaris Volume Manager provided support for high-availability (HA) configurations consisting of two hosts that share at least three strings of drives and that run software enabling exclusive access to the data on those drives from one host. (Note: Volume Manager, by itself, does not actually provide a high-availability environment. The diskset feature is an enabler for HA configurations.)</p> <p>Volume Manager provides support for a low-end HA solution consisting of two hosts that share only two strings of drives. The hosts in this type of configuration, referred to as <i>mediators</i>, run a special daemon, <code>rpc.metamedd(1M)</code>. The mediator hosts take on additional responsibilities to ensure that data is available in the case of host or drive failures.</p> <p>In a mediator configuration, two hosts are physically connected to two strings of drives. This configuration can survive the failure of a single host or a single string of drives, without administrative intervention. If both a host and a string of drives fail (multiple failures), the integrity of the data cannot be guaranteed. At this point, administrative intervention is required to make the data accessible.</p> <p>The following definitions pertain to a mediator configuration:</p> <table><tr><td><code>diskset</code></td><td>A set of drives containing metadevices and hot spares that can be shared exclusively (but not concurrently) by two hosts.</td></tr><tr><td><code>Volume Manager state database</code></td><td>A replicated database that stores metadevice configuration and state information.</td></tr><tr><td><code>mediator host</code></td><td>A host that runs the <code>rpc.metamedd(1M)</code> daemon and that has been added to a diskset. The mediator host participates in checking the state database and the mediator quorum.</td></tr><tr><td><code>mediator quorum</code></td><td>The condition achieved when the number of accessible mediator hosts is equal to half+1 the total number of configured mediator hosts. Because it is expected that there will be two mediator hosts, this number will normally be 2 ($[(2/2) + 1] = 2$).</td></tr><tr><td><code>replica</code></td><td>A single copy of the Volume Manager metadevice state database.</td></tr><tr><td><code>replica quorum</code></td><td>The condition achieved when the number of accessible replicas is equal to half+1 the total number of configured replicas. For example, if a system is configured with ten replicas, the quorum is met when six are accessible ($[(10/2) + 1 = 6]$).</td></tr></table>	<code>diskset</code>	A set of drives containing metadevices and hot spares that can be shared exclusively (but not concurrently) by two hosts.	<code>Volume Manager state database</code>	A replicated database that stores metadevice configuration and state information.	<code>mediator host</code>	A host that runs the <code>rpc.metamedd(1M)</code> daemon and that has been added to a diskset. The mediator host participates in checking the state database and the mediator quorum.	<code>mediator quorum</code>	The condition achieved when the number of accessible mediator hosts is equal to half+1 the total number of configured mediator hosts. Because it is expected that there will be two mediator hosts, this number will normally be 2 ($[(2/2) + 1] = 2$).	<code>replica</code>	A single copy of the Volume Manager metadevice state database.	<code>replica quorum</code>	The condition achieved when the number of accessible replicas is equal to half+1 the total number of configured replicas. For example, if a system is configured with ten replicas, the quorum is met when six are accessible ($[(10/2) + 1 = 6]$).
<code>diskset</code>	A set of drives containing metadevices and hot spares that can be shared exclusively (but not concurrently) by two hosts.												
<code>Volume Manager state database</code>	A replicated database that stores metadevice configuration and state information.												
<code>mediator host</code>	A host that runs the <code>rpc.metamedd(1M)</code> daemon and that has been added to a diskset. The mediator host participates in checking the state database and the mediator quorum.												
<code>mediator quorum</code>	The condition achieved when the number of accessible mediator hosts is equal to half+1 the total number of configured mediator hosts. Because it is expected that there will be two mediator hosts, this number will normally be 2 ($[(2/2) + 1] = 2$).												
<code>replica</code>	A single copy of the Volume Manager metadevice state database.												
<code>replica quorum</code>	The condition achieved when the number of accessible replicas is equal to half+1 the total number of configured replicas. For example, if a system is configured with ten replicas, the quorum is met when six are accessible ($[(10/2) + 1 = 6]$).												

A mediator host running the `rpc.metamedd(1M)` daemon keeps track of replica updates. As long as the following conditions are met, access to data occurs without any administrative intervention:

- The replica quorum is not met.
- Half of the replicas is still accessible.
- The mediator quorum is met.

The following conditions describe the operation of mediator hosts:

1. If the replica quorum is met, access to the diskset is granted. At this point no mediator host is involved.
2. If the replica quorum is not met, half of the replicas is accessible, the mediator quorum is met, and the replica and mediator data match, access to the diskset is granted. The mediator host contributes the deciding vote.
3. If the replica quorum is not met, half of the replicas is accessible, the mediator quorum is not met, half of the mediator hosts is accessible, and the replica and mediator data match, the system prompts you to grant or deny access to the diskset.
4. If the replica quorum is not met, half of the replicas is accessible, the mediator quorum is met, and the replica and mediator data do not match, access to the diskset is read-only. You can delete replicas, release the diskset, and retake the diskset to gain read-write access to the data in the diskset.
5. In all other cases, the diskset access is read-only. You can delete replicas, release the diskset, and retake the diskset to gain read-write access to the data in the diskset.

The `metaset(1M)` command administers disksets and mediator hosts. The following options to the `metaset` command pertain only to administering mediator hosts.

<code>-a -m mediator_host_list</code>	Adds mediator hosts to the named set. A <code>mediator_host_list</code> is the nodename of the mediator host to be added and up to 2 other aliases for the mediator host. The nodename and aliases for each mediator host are separated by commas. Up to 2 mediator hosts can be specified for the named diskset.
<code>-d -m mediator_host_list</code>	Deletes mediator hosts from the named diskset. Mediator hosts are deleted from the diskset by specifying the nodename of mediator host to delete.
<code>-q</code>	Displays an enumerated list of tags pertaining to “tagged data” that may be encountered during a take of the ownership of a diskset.
<code>-t [-f] -y</code>	Takes ownership of a diskset safely, unless <code>-f</code> is used, in which case the take is unconditional. If <code>metaset</code> finds that another host owns the set, this host will not be allowed to take ownership of the set. If the set is not

mediator(7D)

owned by any other host, all the disks within the set will be owned by the host on which `metaset` was executed. The metadvice state database is read in and the shared metadvice devices contained in the set become accessible. The `-t` option will take a diskset that has stale databases. When the databases are stale, `metaset` will exit with code 66, and a message will be printed. At that point, the only operations permitted are the addition and deletion of replicas. Once the addition or deletion of the replicas has been completed, the diskset should be released and retaken to gain full access to the data. If mediator hosts have been configured, some additional exit codes are possible. If half of the replicas and half of the mediator hosts are operating properly, the take will exit with code 3. At this point, you can add or delete replicas, or use the `-y` option on a subsequent take. If the take operation encounters "tagged data," the take operation will exit with code 2. You can then run the `metaset` command with the `-q` option to see an enumerated list of tags.

`-t [-f] -u tagnumber`

Once a tag has been selected, a subsequent take with `-u tagnumber` can be executed to select the data associated with the given `tagnumber`.

SEE ALSO `metaset(1M)`, `md(7D)`, `rpc.metamedd(1M)`, `rpc.metad(1M)`

Sun Cluster documentation, *Solaris Volume Manager Administration Guide*

NOTES Diskset administration, including the addition and deletion of hosts and drives, requires all hosts in the set to be accessible from the network.

NAME	mem, kmem, allkmem – physical or virtual memory access						
SYNOPSIS	<pre> /dev/mem /dev/kmem /dev/allkmem </pre>						
DESCRIPTION	<p>The file <code>/dev/mem</code> is a special file that provides access to the physical memory of the computer.</p> <p>The file <code>/dev/kmem</code> is a special file that provides access to the virtual address space of the operating system kernel, excluding memory that is associated with an I/O device.</p> <p>The file <code>/dev/allkmem</code> is a special file that provides access to the virtual address space of the operating system kernel, including memory that is associated with an I/O device. You can use any of these devices to examine and modify the system.</p> <p>Byte addresses in <code>/dev/mem</code> are interpreted as physical memory addresses. Byte addresses in <code>/dev/kmem</code> and <code>/dev/allkmem</code> are interpreted as kernel virtual memory addresses. A reference to a non-existent location returns an error. See ERRORS for more information.</p> <p>The file <code>/dev/mem</code> accesses physical memory; the size of the file is equal to the amount of physical memory in the computer. This size may be larger than 4GB on a system running the 32-bit operating environment. In this case, you can access memory beyond 4GB using a series of <code>read(2)</code> and <code>write(2)</code> calls, a <code>pread64()</code> or <code>pwrite64()</code> call, or a combination of <code>lseek(2)</code> and <code>read(2)</code> or <code>write(2)</code>.</p>						
ERRORS	<table border="0"> <tr> <td style="vertical-align: top;">EFAULT</td> <td>Occurs when trying to <code>write(2)</code> a read-only location (<code>allkmem</code>), <code>read(2)</code> a write-only location (<code>allkmem</code>), or <code>read(2)</code> or <code>write(2)</code> a non-existent or unimplemented location (<code>mem</code>, <code>kmem</code>, <code>allkmem</code>).</td> </tr> <tr> <td style="vertical-align: top;">EIO</td> <td>Occurs when trying to <code>read(2)</code> or <code>write(2)</code> a memory location that is associated with an I/O device using the <code>/dev/kmem</code> special file.</td> </tr> <tr> <td style="vertical-align: top;">ENXIO</td> <td>Results from attempting to <code>mmap(2)</code> a non-existent physical (<code>mem</code>) or virtual (<code>kmem</code>, <code>allkmem</code>) memory address.</td> </tr> </table>	EFAULT	Occurs when trying to <code>write(2)</code> a read-only location (<code>allkmem</code>), <code>read(2)</code> a write-only location (<code>allkmem</code>), or <code>read(2)</code> or <code>write(2)</code> a non-existent or unimplemented location (<code>mem</code> , <code>kmem</code> , <code>allkmem</code>).	EIO	Occurs when trying to <code>read(2)</code> or <code>write(2)</code> a memory location that is associated with an I/O device using the <code>/dev/kmem</code> special file.	ENXIO	Results from attempting to <code>mmap(2)</code> a non-existent physical (<code>mem</code>) or virtual (<code>kmem</code> , <code>allkmem</code>) memory address.
EFAULT	Occurs when trying to <code>write(2)</code> a read-only location (<code>allkmem</code>), <code>read(2)</code> a write-only location (<code>allkmem</code>), or <code>read(2)</code> or <code>write(2)</code> a non-existent or unimplemented location (<code>mem</code> , <code>kmem</code> , <code>allkmem</code>).						
EIO	Occurs when trying to <code>read(2)</code> or <code>write(2)</code> a memory location that is associated with an I/O device using the <code>/dev/kmem</code> special file.						
ENXIO	Results from attempting to <code>mmap(2)</code> a non-existent physical (<code>mem</code>) or virtual (<code>kmem</code> , <code>allkmem</code>) memory address.						
FILES	<table border="0"> <tr> <td style="vertical-align: top;"><code>/dev/mem</code></td> <td>Provides access to the computer's physical memory.</td> </tr> <tr> <td style="vertical-align: top;"><code>/dev/kmem</code></td> <td>Provides access to the virtual address space of the operating system kernel, excluding memory that is associated with an I/O device.</td> </tr> <tr> <td style="vertical-align: top;"><code>/dev/allkmem</code></td> <td>Provides access to the virtual address space of the operating system kernel, including memory that is associated with an I/O device.</td> </tr> </table>	<code>/dev/mem</code>	Provides access to the computer's physical memory.	<code>/dev/kmem</code>	Provides access to the virtual address space of the operating system kernel, excluding memory that is associated with an I/O device.	<code>/dev/allkmem</code>	Provides access to the virtual address space of the operating system kernel, including memory that is associated with an I/O device.
<code>/dev/mem</code>	Provides access to the computer's physical memory.						
<code>/dev/kmem</code>	Provides access to the virtual address space of the operating system kernel, excluding memory that is associated with an I/O device.						
<code>/dev/allkmem</code>	Provides access to the virtual address space of the operating system kernel, including memory that is associated with an I/O device.						
SEE ALSO	<code>lseek(2)</code> , <code>mmap(2)</code> , <code>read(2)</code> , <code>write(2)</code>						

mem(7D)

WARNINGS Using these devices to modify (that is, write to) the address space of a live running operating system or to modify the state of a hardware device is extremely dangerous and may result in a system panic if kernel data structures are damaged or if device state is changed.

NAME	mhd – multihost disk control operations
SYNOPSIS	<code>#include <sys/mhd.h></code>
DESCRIPTION	<p>The <code>mhd ioctl(2)</code> control access rights of a multihost disk, using disk reservations on the disk device.</p> <p>The stability level of this interface (see <code>attributes(5)</code>) is evolving. As a result, the interface is subject to change and you should limit your use of it.</p> <p>The <code>mhd ioctls</code> fall into two major categories: (1) <code>ioctls</code> for non-shared multihost disks and (2) <code>ioctls</code> for shared multihost disks.</p> <p>One <code>ioctl</code>, <code>MHIOCENFAILFAST</code>, is applicable to both non-shared and shared multihost disks. It is described after the first two categories.</p> <p>All the <code>ioctls</code> require root privilege.</p> <p>For all of the <code>ioctls</code>, the caller should obtain the file descriptor for the device by calling <code>open(2)</code> with the <code>O_NDELAY</code> flag; without the <code>O_NDELAY</code> flag, the <code>open</code> may fail due to another host already having a conflicting reservation on the device. Some of the <code>ioctls</code> below permit the caller to forcibly clear a conflicting reservation held by another host, however, in order to call the <code>ioctl</code>, the caller must first obtain the open file descriptor.</p> <p>Non-shared multihost disks Non-shared multihost disks <code>ioctls</code> consist of <code>MHIOCKDOWN</code>, <code>MHIOCRELEASE</code>, <code>HIOCSTATUS</code>, and <code>MHIOCQRESERVE</code>. These <code>ioctl</code> requests control the access rights of non-shared multihost disks. A non-shared multihost disk is one that supports serialized, mutually exclusive I/O mastery by the connected hosts. This is in contrast to the shared-disk model, in which concurrent access is allowed from more than one host (see below).</p> <p>A non-shared multihost disk can be in one of two states:</p> <ul style="list-style-type: none"> ■ Exclusive access state, where only one connected host has I/O access ■ Non-exclusive access state, where all connected hosts have I/O access. An external hardware reset can cause the disk to enter the non-exclusive access state. <p>Each multihost disk driver views the machine on which it's running as the "local host"; each views all other machines as "remote hosts". For each I/O or <code>ioctl</code> request, the requesting host is the local host.</p> <p>Note that the non-shared <code>ioctls</code> are designed to work with SCSI-2 disks. The SCSI-2 RESERVE/RELEASE command set is the underlying hardware facility in the device that supports the non-shared <code>ioctls</code>.</p> <p>The function prototypes for the non-shared <code>ioctls</code> are:</p>

mhd(7i)

```
ioctl(fd, MHIOCKOWN);  
ioctl(fd, MHIOCRELEASE);  
ioctl(fd, MHIOCSTATUS);  
ioctl(fd, MHIOCQRESERVE);
```

MHIOCKOWN Forcefully acquires exclusive access rights to the multihost disk for the local host. Revokes all access rights to the multihost disk from remote hosts. Causes the disk to enter the exclusive access state.

Implementation Note: Reservations (exclusive access rights) broken via random resets should be reinstated by the driver upon their detection, for example, in the automatic probe function described below.

MHIOCRELEASE Relinquishes exclusive access rights to the multihost disk for the local host. On success, causes the disk to enter the non-exclusive access state.

MHIOCSTATUS Probes a multihost disk to determine whether the local host has access rights to the disk. Returns 0 if the local host has access to the disk, 1 if it doesn't, and -1 with `errno` set to `EIO` if the probe failed for some other reason.

MHIOCQRESERVE Issues, simply and only, a SCSI-2 Reserve command. If the attempt to reserve fails due to the SCSI error Reservation Conflict (which implies that some other host has the device reserved), then the `ioctl` will return -1 with `errno` set to `EACCES`. The `MHIOCQRESERVE` `ioctl` does NOT issue a bus device reset or bus reset prior to attempting the SCSI-2 reserve command. It also does not take care of re-instating reservations that disappear due to bus resets or bus device resets; if that behavior is desired, then the caller can call `MHIOCKOWN` after the `MHIOCQRESERVE` has returned success. If the device does not support the SCSI-2 Reserve command, then the `ioctl` returns -1 with `errno` set to `ENOTSUP`. The `MHIOCQRESERVE` `ioctl` is intended to be used by high-availability or clustering software for a "quorum" disk, hence, the "Q" in the name of the `ioctl`.

Shared Multihost Disks

Shared multihost disks `ioctls` control access to shared multihost disks. The `ioctls` are merely a veneer on the SCSI-3 Persistent Reservation facility. Therefore, the underlying semantic model is not described in detail here, see instead the SCSI-3 standard. The SCSI-3 Persistent Reservations support the concept of a group of hosts all sharing access to a disk.

The function prototypes and descriptions for the shared multihost `ioctls` are as follows:

```
ioctl(fd, MHIOCGRP_INKEYS, (mhioc_inkeys_t) *k);
```

Issues the SCSI-3 command Persistent Reserve In Read Keys to the device. On input, the field `k->li` should be initialized by the caller with `k->li.listsize` reflecting how big of an array the caller has allocated for the `k->li.list` field and

with `k->li.listlen == 0`. On return, the field `k->li.listlen` is updated to indicate the number of reservation keys the device currently has: if this value is larger than `k->li.listsize` then that indicates that the caller should have passed a bigger `k->li.list` array with a bigger `k->li.listsize`. The number of array elements actually written by the callee into `k->li.list` is the minimum of `k->li.listlen` and `k->li.listsize`. The field `k->generation` is updated with the generation information returned by the SCSI-3 Read Keys query. If the device does not support SCSI-3 Persistent Reservations, then this ioctl returns `-1` with `errno` set to `ENOTSUP`.

```
ioctl(fd, MHIOCGRP_INRESVS, (mhioc_inresvs_t) *r);
```

Issues the SCSI-3 command Persistent Reserve In Read Reservations to the device. Remarks similar to `MHIOCGRP_INKEYS` apply to the array manipulation. If the device does not support SCSI-3 Persistent Reservations, then this ioctl returns `-1` with `errno` set to `ENOTSUP`.

```
ioctl(fd, MHIOCGRP_REGISTER, (mhioc_register_t) *r);
```

Issues the SCSI-3 command Persistent Reserve Out Register. The fields of structure `r` are all inputs; none of the fields are modified by the ioctl. The field `r->aptp1` should be set to true to specify that registrations and reservations should persist across device power failures, or to false to specify that registrations and reservations should be cleared upon device power failure; true is the recommended setting. The field `r->oldkey` is the key that the caller believes the device may already have for this host initiator; if the caller believes that that this host initiator is not already registered with this device, it should pass the special key of all zeros. To achieve the effect of unregistering with the device, the caller should pass its current key for the `r->oldkey` field and an `r->newkey` field containing the special key of all zeros. If the device returns the SCSI error code Reservation Conflict, this ioctl returns `-1` with `errno` set to `EACCES`.

```
ioctl(fd, MHIOCGRP_RESERVE, (mhioc_resv_desc_t) *r);
```

Issues the SCSI-3 command Persistent Reserve Out Reserve. The fields of structure `r` are all inputs; none of the fields are modified by the ioctl. If the device returns the SCSI error code Reservation Conflict, this ioctl returns `-1` with `errno` set to `EACCES`.

```
ioctl(fd, MHIOCGRP_PREEMPTANDABORT, (mhioc_preemptandabort_t) *r);
```

Issues the SCSI-3 command Persistent Reserve Out Preempt-And-Abort. The fields of structure `r` are all inputs; none of the fields are modified by the ioctl. The key of the victim host is specified by the field `r->victim_key`. The field `r->resvdesc` supplies the preempter's key and the reservation that it is requesting as part of the SCSI-3 Preempt-And-Abort command. If the device returns the SCSI error code Reservation Conflict, this ioctl returns `-1` with `errno` set to `EACCES`.

```
ioctl(fd, MHIOCGRP_PREEMPT, (mhioc_preemptandabort_t) *r);
```

Similar to `MHIOCGRP_PREEMPTANDABORT`, but instead issues the SCSI-3 command Persistent Reserve Out Preempt.

```
ioctl(fd, MHIOCGRP_CLEAR, (mhioc_resv_key_t) *r);
```

Issues the SCSI-3 command Persistent Reserve Out Clear. The input parameter `r` is the reservation key of the caller, which should have been already registered with

mhd(7i)

the device, by an earlier call to `MHIOCGRP_REGISTER`.

For each device, the non-shared ioctls should not be mixed with the Persistent Reserve Out shared ioctls, and vice-versa, otherwise, the underlying device is likely to return errors, because SCSI does not permit SCSI-2 reservations to be mixed with SCSI-3 reservations on a single device. It is, however, legitimate to call the Persistent Reserve In ioctls, because these are query only. Issuing the `MHIOCGRP_INKEYS` ioctl is the recommended way for a caller to determine if the device supports SCSI-3 Persistent Reservations (the ioctl will return `-1` with `errno` set to `ENOTSUP` if the device does not).

MHIOCENFAILFAST The `MHIOCENFAILFAST` ioctl is applicable for both non-shared and shared disks, and may be used with either the non-shared or shared ioctls.

`ioctl(fd, MHIOCENFAILFAST, (unsigned int *) millisecs) ;`

Enables or disables the failfast option in the multihost disk driver and enables or disables automatic probing of a multihost disk, described below. The argument is an unsigned integer specifying the number of milliseconds to wait between executions of the automatic probe function. An argument of zero disables the failfast option and disables automatic probing. If the `MHIOCENFAILFAST` ioctl is never called, the effect is defined to be that both the failfast option and automatic probing are disabled.

Automatic Probing

The `MHIOCENFAILFAST` ioctl sets up a timeout in the driver to periodically schedule automatic probes of the disk. The automatic probe function works in this manner: The driver is scheduled to probe the multihost disk every `n` milliseconds, rounded up to the next integral multiple of the system clock's resolution. If

1. the local host no longer has access rights to the multihost disk, and
2. access rights were expected to be held by the local host,

the driver immediately panics the machine to comply with the failfast model.

If the driver makes this discovery outside the timeout function, especially during a read or write operation, it is imperative that it panic the system then as well.

RETURN VALUES

Each request returns `-1` on failure and sets `errno` to indicate the error.

`EPERM` Caller is not root.

`EACCESS` Access rights were denied.

`EIO` The multihost disk or controller was unable to successfully complete the requested operation.

`EOPNOTSUP` The multihost disk does not support the operation. For example, it does not support the SCSI-2 Reserve/Release command set, or the SCSI-3 Persistent Reservation command set.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWhea
Stability	Evolving

SEE ALSO `ioctl(2)`, `open(2)`, `attributes(5)`, `open(2)`

mixer(7I)

NAME	mixer – audio mixer audio personality module interface
SYNOPSIS	<pre>#include <sys/mixer.h></pre>
OVERVIEW	The <code>audio mixer</code> extends the <code>audio(7I)</code> interface, allowing more than one process to play or record audio at the same time. Understanding the <code>audio(7I)</code> interface thoroughly is a prerequisite to understanding the <code>mixer(7I)</code> interface.
Backward Compatibility	It is possible to disable the mixing function and return to 100% backward compatibility with the <code>audio(7I)</code> interface. These two modes of operation are referred to as the mixer mode and the compatible mode. This is done by using either the <code>mixerctl(1)</code> or <code>sdtaudiocontrol(1)</code> applications, or by editing the audio driver's <code>.conf</code> file and then unloading and reloading the audio driver.
Multi-Stream Codecs	The audio mixer supports multi-stream Codecs. Examples of these Codecs are the Crystal Semiconductor 4410/4422 and the Aureal 8820/8830. These devices have DSP engines that provide sample rate conversion and other features. Each play/record channel is mapped to an individual channel straight into the Codec. The audio mixer does not perform sample rate or encoding conversion. (See below). However, the programming interfaces remain the same and applications cannot distinguish between multi-stream Codec and traditional Codec.
Buffer Size	<p>An application can use the audio information structure to set the size of the play/record buffers. As with the <code>audio(7I)</code> interface, the audio mixer does not support changing of the play buffer size. Instead, the audio driver takes sound samples as it can handle them, regardless of how many are delivered to the driver with each write.</p> <p>The audio mixer supports changing of the record buffer size. When captured by the audio driver, <i>buffer size</i> bytes are sent to the application for reading.</p>
AUDIO FORMATS	See the <code>audio(7I)</code> manual page for a brief discussion of audio formats. To mix the various audio streams, the audio mixer must convert all audio formats to a common format. The following describes how the audio mixer deals with these different components.
Sample Rate	<p>When <code>/dev/audio</code> is opened, the initial sample rate is 8KHz, as defined in <code>audio(7I)</code>.</p> <p>In mixer mode, the audio mixer always configures the Codec for the highest possible sample rate for both play and record. This ensures that none of the audio streams require compute-intensive low pass filtering. The result is that high sample rate audio streams are not degraded by filtering.</p> <p>Sample rate conversion can be a compute-intensive operation, depending on the number of channels and a device's sample rate. For example, an 8KHz signal can be easily converted to 48KHz, requiring a low cost up sampling by 6. However, converting from 44.1KHz to 48KHz <i>is</i> compute intensive because it must be up sampled by 160 and then down sampled by 147. (This is only done using integer multipliers.)</p>

Encodings and Precision

Applications can greatly reduce the impact of sample rate conversion by carefully picking the sample rate. Applications should always use the highest sample rate the device supports. An application can also do its own sample rate conversion (to take advantage of floating point and accelerated instruction) or use small integers for up and down sampling.

In compatible mode, the audio mixer programs the Codec to the sample rate set by the application to avoid incurring any sample rate conversion overhead. If the Codec cannot support different play and record sample rates, the `AUDIO_SETINFO ioctl(2)` fails.

When `/dev/audio` is opened, initial encoding and precision is 8-bit μ -Law (as in the Greek letter mu) . (As defined in [audio\(7I\)](#))

In mixer mode, the audio mixer supports formats in the following precisions:

Encoding	Precision	Channels
Signed Linear PCM	16-bit	Mono or Stereo
Signed Linear PCM	8-bit	Mono or Stereo
μ -Law	8-bit	Mono or Stereo
A-Law	8-bit	Mono or Stereo

The audio mixer converts all audio streams to 16-bit Linear PCM before mixing. After mixing, conversion is made to the best possible Codec format. The conversion process is not compute intensive and audio applications can choose the encoding format that best meets its needs.

In compatibility mode, the audio mixer sets the Codec to the encoding and precision set by the application. If the Codec cannot support different play and record encodings or precisions, the `AUDIO_SETINFO ioctl(2)` call fails.

Channels

When `/dev/audio` is opened, the number of initial channels is 1, left channel mono. (As defined in [audio\(7I\)](#)). Most Codecs play or record mono audio on the left channel.

In mixer mode, the audio mixer sets the Codec to the maximum number of channels supported. If a mono signal is played or recorded, it is mixed on the first (usually the left) channel only. Silence is mixed on the remaining channels.

In compatible mode, the audio mixer sets the Codec to the number of channels set by the application. If the Codec cannot support a different number of play and record channels, the `AUDIO_SETINFO ioctl(2)` call fails.

mixer(7I)

DESCRIPTION

The device `/dev/audio` is a device driver that dispatches audio requests to the appropriate underlying audio personality module. The audio driver is implemented as a STREAMS driver. To record audio input, applications `open(2)` the `/dev/audio` device and read data from it using the `read(2)` system call. Similarly, sound data is queued to the audio output port by using the `write(2)` system call. Device configuration is performed using the `ioctl(2)` interface.

Opening the Audio Device

In mixer mode, the audio device is no longer treated as an exclusive resource. Each process may open the audio device once unless the process has made an `AUDIO_MIXER_MULTIPLE_OPEN` `ioctl(2)` call. See below for details.

Each `open()` will complete as long as there are channels available to be allocated. If no channels are available to be allocated:

- if either the `O_NDELAY` or `O_NONBLOCK` flags are set in the `open()` *oflag* argument, `-1` is immediately returned, with *errno* set to `EBUSY`.
- if neither the `O_NDELAY` nor the `O_NONBLOCK` flags are set, then `open()` hangs until a channel becomes available or a signal is delivered to the process. In the latter case, a `-1` is returned with *errno* set to `EINTR`.

Upon the initial `open()` of the audio channel, the audio mixer sets the data format of the audio channel to the default state of 8-bit, 8Khz, mono μ -Law data. If the audio device does not support this configuration, it informs the audio mixer of the initial configuration. Audio applications should explicitly set the encoding characteristics to match the audio data requirements, and not depend on the default configuration. See the [audio\(7I\)](#) manual page for details on how the audio mixer behaves when in compatible mode.

Recording Audio Data

The `read(2)` system call copies data from the system buffers to the application. Ordinarily, `read()` blocks until the user buffer is filled. The `I_NREAD` `ioctl` (see [streamio\(7I\)](#)) may be used to determine the amount of data that may be read without blocking. The device may also be set to a non-blocking mode, where `read()` completes immediately but may return fewer bytes than requested. See the `read(2)` manual page for a complete description of this behavior.

When the audio device is opened with read access, the device driver immediately starts buffering audio input data. Because this consumes system resources, processes that do not record audio data should open the device write-only (`O_WRONLY`).

The transfer of input data to STREAMS buffers may be paused (or resumed) by using the `AUDIO_SETINFO` `ioctl` to set (or clear) the `record.pause` flag in the audio information structure. (See [audio\(7I\)](#).) All unread input data in the STREAMS queue may be discarded by using the `I_FLUSH_STREAMS` `ioctl` (see [streamio\(7I\)](#)). When changing record parameters, the input stream should first be paused and flushed before the change. Otherwise, subsequent reads may return samples in the old format, followed by samples in the new format.

Input data accumulates in STREAMS buffers rapidly. For 8-bit, 8 KHz, mono μ -Law data, it accumulates at 8000 bytes per second. If a device is configured for 16-bit linear or higher sample rates, it accumulates even faster. If the application that consumes the

	<p>data is unable to meet the input data rate, the STREAMS queue may become full. When this happens, the <code>record.error</code> flag is set in the audio information structure and input sampling ceases until there is room for additional data, resulting in a data stream discontinuity. To prevent this, audio recording applications should open the audio device when they are ready to begin reading data and not at the start of extensive initialization.</p>
Playing Audio Data	<p>The <code>write(2)</code> system call copies data from an application's buffer to the STREAMS output queue. Ordinarily, <code>write()</code> blocks until the entire user buffer is transferred. The device may alternatively be set to a non-blocking mode, in which case <code>write()</code> completes immediately, but may transfer fewer bytes than requested. (See the <code>write(2)</code> manual page for a complete description of this behavior).</p> <p>Although <code>write()</code> returns when the data is successfully queued, the actual completion of audio output may take considerably longer. The <code>AUDIO_DRAIN</code> <code>ioctl</code> may be issued to allow an application to block until all of the queued output data has been played. Alternatively, a process may request asynchronous notification of output completion by writing a zero-length buffer (end-of-file record) to the output stream. When such a buffer has been processed, the <code>play.eof</code> flag in the audio information structure (see below) is incremented.</p> <p>The final <code>close(2)</code> of the audio device file descriptor hangs until all of the process' remaining audio output has drained. If a signal interrupts the <code>close()</code> or if the process exits without closing the audio device, any remaining data queued for audio output is flushed and the audio device is closed immediately.</p> <p>The conversion of output data may be paused (or resumed) by using the <code>AUDIO_SETINFO</code> <code>ioctl</code> to set (or clear) the <code>play.pause</code> flag in the audio information structure. Queued output data may be discarded by using the <code>I_FLUSH</code> STREAMS <code>ioctl</code>. (See streamio(7I).)</p> <p>Output data is played from the STREAMS buffers at a default rate of 8000 bytes per second for μ-Law, A-Law, or 8-bit PCM data, or at a faster rate for 16-bit linear data or higher sampling rates. If the output queue becomes empty, the <code>play.error</code> flag is set in the audio information structure and output is stopped until additional data is queued. If an application attempts to write a number of bytes that is not a multiple of the current sample frame size, an error is generated and the bad data is thrown away. Additional writes are allowed.</p>
Asynchronous I/O	<p>The <code>I_SETSIG</code> STREAMS <code>ioctl</code> (see streamio(7I)) enables asynchronous notification through the <code>SIGPOLL</code> signal of input and output ready conditions. The <code>O_NONBLOCK</code> flag may be set using the <code>F_SETFL</code> <code>fcntl(2)</code> to enable non-blocking <code>read()</code> and <code>write()</code> requests. This is normally sufficient for applications to maintain a background audio stream.</p>
Audio Control Pseudo-Device	<p>The <code>/dev/audioctl</code> pseudo-device enables an application to modify characteristics of the audio device while it is being used by an unrelated process. Any number of processes may open the <code>/dev/audioctl</code> pseudo device simultaneously. <code>/dev/audioctl</code> ignores <code>read()</code> and <code>write()</code> system calls.</p>

mixer(7I)

Note – The audio control device name is constructed by appending the letters "ctl" to the path name of the audio device.

Audio Status Change Notification

Applications that open the audio control pseudo-device may request asynchronous notification of changes in the state of the audio device by setting the `S_MSG` flag in an `I_SETSIG STREAMS` `ioctl`. (See [streamio\(7I\)](#)). Such processes receive a `SIGPOLL` signal when any of the following events occur:

- `AUDIO_SETINFO`, `AUDIO_MIXERCTL_SETINFO`, `AUDIO_MIXERCTL_SET_CHINFO`, or `AUDIO_MIXERCTL_SET_MODE` `ioctl` () has altered the device state.
- Input overflow or output underflow has occurred.
- End-of-file record (zero-length buffer) has been processed on output.
- `open()` or `close()` of `/dev/audio` has altered the device state.
- An external event (such as speakerbox's volume control) has altered the device state.

IOCTLS

The audio mixer implements all the `ioctl` ()s defined in [audio\(7I\)](#) and uses the `audio_prinfo_t`, `audio_info_t`, and `audio_device_t` structures. See the [audio\(7I\)](#) manual page for details on these `ioctl` ()s and structures. The audio mixer also uses the data structures described below.

Audio Mixer Control Structure

The state of the audio device may be polled or modified using the `AUDIO_MIXERCTL_GETINFO` and `AUDIO_MIXERCTL_SETINFO` `ioctl` commands.

```
struct am_control {
    audio_info_t dev_info;      /* the audio device's state */
    int8_t      ch_open[1];    /* variable sized array of */
                                /* of open channels */
};
typedef struct am_control_t;
```

See `CODE EXAMPLES` for sample code on how to use this structure and the related macro `AUDIO_MIXER_CTL_STRUCT_SIZE(num_ch)`.

Audio Mixer Sample Rates Structure

The following structure is used by the `AUDIO_MIXER_GET_SAMPLE_RATES` `ioctl` to get a list of all the supported sample rates.

```
struct am_sample_rates {

    uint_t type;              /* play or capture */
    uint_t flags;
    uint_t num_samp_rates;    /* number of elements */
    /* in samp_rates[] */
    uint_t samp_rates[1];    /*variable sized array */
    /* of sample rates */
};
typedef struct am_sample_rates am_sample_rates_t;

#define AUDIO_PLAY 1 /* output */
#define AUDIO_RECORD 2 /* input */
```


	<pre>#define MIXER_SR_LIMITS 0x00000001/* sample rates */ /* set limits */</pre>
	<p>See CODE EXAMPLES for example code on how to use this structure and the related macro <code>AUDIO_MIXER_SAMP_RATES_STRUCT_SIZE(num_srs)</code>.</p>
Audio Info Structure	<p>When in mixer mode, the <code>AM_MIXER</code> bit in the <code>audio_info_t</code> structure's <code>sw_features_enabled</code> field is set. When in compatibility mode, that bit is clear.</p> <p>The defines for the <code>sw_features</code> and the <code>sw_features_enabled</code> fields are:</p> <pre>#define AM_MIXER 0x00000001 /* mixer is present/enabled */</pre>
Streamio IOCTLS	<p>All <code>streamio(7I)</code> <code>ioctl</code> commands may be issued for the <code>/dev/audio</code> and <code>/dev/audiocpl</code> devices. <code>I_SETSIG</code> <code>ioctl</code> may be issued for <code>/dev/audiocpl</code> to enable the notification of audio status changes as described above.</p>
Audio Mixer IOCTLS	<p>Except for <code>AUDIO_MIXER_GET_SAMPLE_RATE</code>, <code>AUDIO_MIXERCTL_GET_MODE</code>, and <code>AUDIO_MIXERCTL_SET_MODE</code>, these <code>ioctl</code> ()s are valid only in mixer mode. Using them in compatible mode returns an <code>EINVAL</code> error.</p> <p>AUDIO_MIXER_MULTIPLE_OPEN This command allows an individual process to open <code>/dev/audio</code> more than once for play or record. The argument is ignored. This feature is useful for mixing panels that may control multiple audio streams.</p> <p>AUDIO_MIXER_SINGLE_OPEN This command gets a list of supported sample rates for either play or record for the audio mixer's current mode. The argument is ignored. This command returns <code>/dev/audio</code> back to an exclusive access device on per process basis after an <code>AUDIO_MIXER_MULTIPLE_OPEN</code> <code>ioctl</code> () is executed. This <code>ioctl</code> () fails if the process has more than one play or record stream open.</p> <p>AUDIO_MIXER_GET_SAMPLE_RATES This command retrieves sample rates. The argument is a pointer to an <code>am_samples_rates_t</code> structure. It is legal for the supported sample rates to be different for mixer mode vs. compatible mode. The <code>type</code> field must be set to either <code>AUDIO_PLAY</code> or <code>AUDIO_RECORD</code> to get a list of either play or record sample rates, respectively. Setting to both or neither is an error. The <code>num_samp_rates</code> field is set to the number of sample rates that the <code>samp_rates[]</code> array may hold. When the <code>ioctl</code> returns, <code>num_samp_rates</code> is set either to the number of sample rates in the array <code>samp_rates[]</code>, or the total number of sample rates available if there are more than the array can hold. In the former case, there are <code>num_samp_rates</code> valid sample rates in the array. In the latter case, all the elements of the array have valid sample rates, but there are more available. The size of the array should be increased to get all available sample rates. If the <code>flags</code> field has the <code>MIXER_SR_LIMITS</code> flag set, the returned sample rates are the lowest and the highest rates possible, with all sample rates in-between being legal. Some Codecs that have DSP engines on them have this capability.</p> <p>AUDIO_MIXERCTL_GETINFO This command gets device and channel state information. The argument is a pointer to an <code>am_control_t</code> structure. The <code>dev_info</code> field contains the state of</p>

the hardware device. It provides a convenient way to determine the hardware's state. The `ch_open` array is used to specify which channels are open and which are closed. Open channels have non-zero values, closed channels are set to zero, The channel number corresponds to the array index. The number of elements in the `ch_open` array may change over time and a macro is provided to allocate the correct amount of space. The MACROS section below provides more information.

AUDIO_MIXERCTL_SETINFO

This command sets the device state but cannot modify any channel's state. (Use `AUDIO_MIXERCTL_SET_CHINFO` (below) to modify a channel's state.) The argument is a pointer to an `am_control_t` structure. The `dev_info` field sets the device state. The `dev_info` field is used to set the device state. However, there are several limitations. Only the *gain*, *balance*, *port* and *pause* fields for play and record, *monitor_gain* and *output_muted* fields may be modified. (Modifying other fields would interfere with how the audio mixer programs the audio device.) The `ch_open` array is not used when setting the audio device and may be set to a size of one.

AUDIO_MIXERCTL_GET_CHINFO

This command gets a channel's state information. The argument is a pointer to an `audio_channel_t` structure. This command gets a channel's state information. To enable the audio mixer to determine channel information, set the `ch_number` field before making the `ioctl()` call. The `info_size` field must be set to the size of the `audio_info_t` structure. The `*info` field must point to an `audio_info_t` structure. When the `ioctl()` returns, the `pid` field should be checked. If it is set to 0, the remaining data in the `audio_channel_t` structure is invalid because the channel has not been allocated. The `dev_type` field describes the type of channel; the `*info` pointer points to a buffer where the `audio_info_t` structure for the audio channel is populated.

AUDIO_MIXERCTL_SET_CHINFO

This command sets a channel's state information. The argument is a pointer to an `audio_channel_t` structure. Prior to issuing the `ioctl` call, specify the channel to be set in the argument's `ch_number` field, set `*info` to point to an `audio_info_t` structure containing info used to program the state of the channel, and set the `info_size` field to the size of an `audio_info_t` structure. When the `ioctl()` returns, the `pid` field contains the process ID of the process that has the channel open, and `dev_type` contains the type of the device. If `pid` is 0 (zero), the channel is not open.

AUDIO_MIXERCTL_GET_MODE

This command retrieves the mode of the audio mixer. The argument is a pointer to an integer that contains the audio mixer's mode upon return. It is set to either `AM_MIXER_MODE` for mixer mode or `AM_COMPAT_MODE` for compatible mode.

AUDIO_MIXERCTL_SET_MODE

This command sets the mode of the audio mixer. The argument is a pointer to an integer that contains the audio mixer mode to be set. It must be set to either `AM_MIXER_MODE` or `AM_COMPAT_MODE`. The audio mixer may be set to mixer mode at any time, but can only be set to compatible mode when there is a single

read/write open within one process, or a single read process and a single write process. Otherwise the `ioctl()` will fail. Because the Codec is being reprogrammed to a different data format, there may be brief pause or burst of noise when the mode changes. This can be eliminated by pausing the input and output or by closing all streams before changing modes. The `sdaudiocontrol(1)` or `mixerctl(1)` commands may be used to change the audio mixer's mode.

MACROS The following macro is used to determine how large an `am_control_t` structure is when it points to an `audio_info_t` structure.

```
AUDIO_MIXER_CTL_STRUCT_SIZE(num_ch)
```

Where `num_ch` is the number of channels the device supports. The number of channels can be determined using the `AUDIO_GET_NUM_CHS` `ioctl()`.

This macro is used when allocating an `am_sample_rates_t` structure.

```
AUDIO_MIXER_SAMP_RATES_STRUCT_SIZE(num_srs)
```

Where `num_srs` is the number of samples rates requested.

CODE EXAMPLES The following examples illustrate how these new data structures and `ioctls` can be used.

Example 1 The following code demonstrates how to use the audio support and the audio mixer `ioctl()`s to get state information on `/dev/audio`.

```
audio_channel_t ch;
audio_info_t    info;
am_control_t    *ctl;
int             num;

err = ioctl(audio_fd, AUDIO_GET_NUM_CHS, &num);

ctl = (am_control_t *)malloc(AUDIO_MIXER_CTL_STRUCT_SIZE(num));

err = ioctl(audio_fd, AUDIO_MIXERCTL_GETINFO, ctl);

ch.info = &info;
ch.info_size = sizeof (audio_info_t);

for (i = 0; i < num; i++) {
    if (ctl->ch_open[i] != 0) {
        ch.ch_number = i;
        if (ioctl(audio_fd, AUDIO_MIXERCTL_GET_CHINFO, &ch) < 0) {
            printf("Channel # %d isn't an audio/audioctl device\n", i);
        } else {
            printf("Ch# %d, PID = %d, Type = %d\n",
                i, ch.pid, ch.dev_type);
        }
    }
}
}
```

mixer(7I)

Example 2 | The following code demonstrates how to use the `AUDIO_MIXER_GET_SAMPLE_RATES` `ioctl` to get the number of supported play sample rates. It also shows how to deal with allocating a `samp_rates[]` array that is too small.

```
#define LARGE_NUMBER 10000;
am_sample_rates_t *sr;
int num;

for (num = 4; num < LARGE_NUMBER; num += 2) {
    sr = (am_sample_rates_t *)
        malloc(AUDIO_MIXER_SAMP_RATES_STRUCT_SIZE(num));

    sr->num_samp_rates = num;
    sr->type = AUDIO_PLAY;

    err = ioctl(audio_fd, AUDIO_MIXER_GET_SAMPLE_RATES, sr);

    if (sr->num_samp_rates <= num) {
        break;
    }
    free(sr);
}
(void) printf("Supported play sample rates:\n");
for (i = 0; i < sr->num_samp_rates; i++) {
    (void) printf(" %d\n", sr->samp_rates[i]);
}
```

ERRORS | An `open()` fails if:

EBUSY	The requested play or record access is busy and either the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flag was set in the <code>open()</code> request.
ENOMEM	Memory was not available to be allocated for the channel.
EINTR	The requested play or record access is busy and a signal interrupted the <code>open()</code> request.
EIO	There has been an error opening the device. An error message is printed on the console explaining the failure.

An `ioctl()` will fail if:

EBUSY	The parameter changes requested in the <code>AUDIO_SETINFO</code> <code>ioctl</code> could not be made because another process has the device open and is using a different format.
EINTR	The <code>ioctl()</code> was interrupted by a signal.
EINVAL	The parameter changes requested in the <code>AUDIO_SETINFO</code> <code>ioctl</code> are invalid or are not supported by the device, or the audio mixer is in compatible mode and a mixer-mode-only audio mixer <code>ioctl</code> was issued.
EIO	There has been an error with the <code>ioctl()</code> . An error message is printed on the console explaining the failure.

ENOMEM The `ioctl()` failed because memory couldn't be allocated.

FILES The physical audio device names are system dependent and are rarely used by programmers. The programmer should use the generic device names listed below.

- `/dev/audio` symbolic link to the system's primary audio device
- `/dev/audioctrl` symbolic link to the control device for `/dev/audio`
- `/dev/sound/0` first audio device in the system
- `/dev/sound/0ctrl` audio control device for `/dev/sound/0`
- `/dev/sound/x` additional audio devices
- `/dev/sound/xctrl` audio control device for `/dev/sound/x`

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86
Availability	SUNWaudd, SUNWauddx, SUNWaudh, SUNWauda
Stability Level	Evolving

SEE ALSO `mixerctl(1)`, `sdaudiocontrol(1)`, `close(2)`, `fcntl(2)`, `ioctl(2)`, `open(2)`, `poll(2)`, `read(2)`, `write(2)`, `attributes(5)`, `audiocs(7D)`, `audioens(7D)`, `audiots(7D)`, `usb_ac(7D)`, `audio(7I)`, `audio_support(7I)`, `streamio(7I)`

BUGS Due to a feature of the STREAMS implementation, programs that are terminated or exit without closing the audio device may hang for a short period while audio output drains. In general, programs that produce audio output should catch the `SIGINT` signal and flush the output stream before exiting.

mpt(7D)

NAME	mpt – SCSI host bus adapter driver
SYNOPSIS	<code>scsi@unit-address</code>
DESCRIPTION	<p>The <code>mpt</code> host bus adapter driver is a SCSI compliant nexus driver that supports the LSI 53C1030 SCSI and SAS1064 chips.</p> <p>The <code>mpt</code> driver supports the standard functions provided by the SCSI interface, including tagged and untagged queuing, Narrow/Wide/Fast/Ultra SCSI/Ultra SCSI 2/Ultra SCSI 3/Ultra SCSI 4, and auto request sense. The <code>mpt</code> driver does not support linked commands.</p>
DRIVER CONFIGURATION	<p>You configure the <code>mpt</code> driver by defining properties in <code>mpt.conf</code>. These properties override the global SCSI settings. The <code>mpt</code> driver supports these modifiable properties: <code>scsi-options</code>, <code>target<n>-scsi-options</code>, <code>scsi-reset-delay</code>, <code>scsi-tag-age-limit</code>, <code>scsi-watchdog-tick</code>, and <code>scsi-initiator-id</code>.</p> <p>The property <code>target<n>-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target<n></code>, where <code><n></code> can vary from decimal 0 to 15. The <code>mpt</code> driver supports the following SCSI options: <code>SCSI_OPTIONS_DR</code>, <code>SCSI_OPTIONS_SYNC</code>, <code>SCSI_OPTIONS_TAG</code>, <code>SCSI_OPTIONS_FAST</code>, <code>SCSI_OPTIONS_WIDE</code>, <code>SCSI_OPTIONS_FAST20</code>, <code>SCSI_OPTIONS_FAST40</code>, <code>SCSI_OPTIONS_FAST80</code>, <code>SCSI_OPTIONS_FAST160</code>, and <code>SCSI_OPTIONS_QAS</code>.</p> <p>After periodic interval <code>scsi-watchdog-tick</code>, the <code>mpt</code> driver searches through all current and disconnected commands for timeouts.</p> <p>The <code>scsi-tag-age-limit</code> property represents the number of times that the <code>mpt</code> driver attempts to allocate a tag ID that is currently in use after going through all tag IDs in a circular fashion. When encountering the same tag ID used <code>scsi-tag-age-limit</code> times, no additional commands are submitted to the target until all outstanding commands complete or timeout.</p> <p>Refer to <code>scsi_hba_attach(9F)</code> for details.</p>
EXAMPLES	<p>EXAMPLE 1 Example 1: Using the <code>mpt</code> Configuration File</p> <p>Create a file called <code>/kernel/drv/mpt.conf</code>, then add the following line:</p> <pre>scsi-options=0x78;</pre> <p>The above example disables tagged queuing, Fast/Ultra SCSI, and wide mode for all <code>mpt</code> instances.</p> <p>The following example disables an option for one specific <code>mpt</code> device. See <code>driver.conf(4)</code> and <code>pci(4)</code> for more details.</p> <pre>name="mpt" parent="/pci@1f,4000" unit-address="3" target1-scsi-options=0x58 scsi-options=0x178 scsi-initiator-id=6;</pre> <p>Note that the default initiator ID in OBP is 7 and that the change to ID 6 occurs at attach time. It may be preferable to change the initiator ID in OBP.</p>

EXAMPLE 1 Example 1: Using the mpt Configuration File (Continued)

The example above sets `scsi-options` for target 1 to 0x58 and all other targets on this SCSI bus to 0x178.

You can determine the physical path name of the parent by using the `/devices` tree or by following the link of the logical device name:

```
# ls -l /dev/rdisk/c0t0d0s0
lrwxrwxrwx 1 root root 45 May 16 10:08 /dev/rdisk/c0t0d0s0 ->
. . / . . /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

As in the previous example, the parent is `/pci@1f,4000` and the unit-address is the number bound to the `scsi@3` node.

To set `scsi-options` more specifically per target, do the following:

```
target1-scsi-options=0x78;
device-type-scsi-options-list =
    "SEAGATE ST32550W", "seagate-scsi-options" ;
seagate-scsi-options = 0x58;
scsi-options=0x3f8;
```

The above sets `scsi-options` for target 1 to 0x78. All other targets on the SCSI bus are set to 0x3f8 (with the exception of one specific disk type for which `scsi-options` is set to 0x58).

`scsi-options` specified per target ID have the highest precedence, followed by `scsi-options` per device type. Global `scsi-options` (for all mpt instances) per bus have the lowest precedence.

You must reboot the system for the specified `scsi-options` to take effect.

Driver Capabilities

To enable some driver features, the target driver must set capabilities in the mpt driver. The following capabilities can be queried and modified by the target driver: `synchronous`, `tagged-qing`, `wide-xfer`, `auto-rqsense`, `qfull-retries`, and `qfull-retry-interval`. All other capabilities are query only.

By default, `tagged-qing`, `auto-rqsense`, and `wide-xfer` capabilities are disabled, while `disconnect`, `synchronous`, and `untagged-qing` are enabled. These capabilities can only have binary values (0 or 1). The default value for `qfull-retries` is 10, while the default value for `qfull-retry-interval` is 100. The `qfull-retries` capability is a `uchar_t` (0 to 255), while `qfull-retry-interval` is a `ushort_t` (0 to 65535).

The target driver must enable `tagged-qing` and `wide-xfer` explicitly. The `untagged-qing` capability is always enabled and its value cannot be modified.

If a conflict exists between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only whom `!= 0` is supported in the `scsi_ifsetcap(9F)` call.

Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

mpt(7D)

FILES /kernel/drv/mpt 32-bit ELF kernel module
/kernel/drv/sparcv9/mpt64-bit SPARC ELF kernel module
/kernel/drv/amd64/mpt 64-bit x86 ELF kernel module
/kernel/drv/mpt.conf optional configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86, SPARC (see below)

This driver is limited to PCI-based systems with LSI 53c1030 SCSI I/O processors.

SEE ALSO `prtconf(1M)`, `driver.conf(4)`, `pci(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

ANSI Small Computer System Interface-2 (SCSI-2)

SYM53c1030 PCI-SCSI I/O processor Dual Channel Fast-160 — LSI Logic Inc.

DIAGNOSTICS The messages described below are logged and may also appear on the system console.

Device is using a hilevel intr

The device was configured with an interrupt level that cannot be used with this mtp driver. Check the PCI device.

Map setup failed

The driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices are inaccessible.

Cannot map configuration space

The driver was unable to map in the configuration registers. Check for bad hardware. SCSI devices will be inaccessible.

Attach failed

The driver was unable to attach; usually preceded by another warning that indicates why attach failed. These can be considered hardware failures.

Connected command timeout for Target <id>.

This is usually a SCSI bus problem. Check cables and termination.

Target <id> reducing sync. transfer rate

A data transfer hang or DATA-IN phase parity error was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id> reverting to async. mode

A second data transfer hang was detected for this target. The driver attempts to eliminate this problem by reducing the data transfer rate.

Target <id> disabled wide SCSI mode

A second data phase hang was detected for this target. The driver attempts to eliminate this problem by disabling wide SCSI mode.

NOTES The mpt driver supports LSI 53c1030 chips. The LSI 53c1030 chip series support Wide, Fast and Ultra SCSI 4 mode. The maximum LVD SCSI bandwidth is 320 MB/sec.

The mpt driver exports properties indicating the negotiated transfer speed per target (target<n>-sync-speed), whether wide bus is supported (target<n>-wide) for that particular target (target<n>-scsi-options), and whether tagged queuing has been enabled (target<n>-TQ). The sync-speed property value indicates the data transfer rate in KB/sec. The target<n>-TQ and the target<n>-wide property have value 1 (to indicate that the corresponding capability is enabled for that target), or 0 (to indicate that the capability is disabled for that target). See prtconf(1M) (verbose option) for details on viewing the mpt properties.

```
scsi, instance #4
  Driver properties:
    name='target8-TQ' type=int items=1 dev=none
      value=00000001
    name='target8-wide' type=int items=1 dev=none
      value=00000001
    name='target8-sync-speed' type=int items=1 dev=none
      value=00013880
    name='target5-TQ' type=int items=1 dev=none
      value=00000001
    name='target5-wide' type=int items=1 dev=none
      value=00000001
    name='target5-sync-speed' type=int items=1 dev=none
      value=00013880
    name='target4-TQ' type=int items=1 dev=none
      value=00000001
    name='target4-wide' type=int items=1 dev=none
      value=00000001
    name='target4-sync-speed' type=int items=1 dev=none
      value=00013880
    name='pm-components' type=string items=3 dev=none
      value='NAME=mpt4' + '0=Off (PCI D3 State)' + '3=On (PCI \
        DO State)'
```

msglog(7D)

- NAME** msglog – message output collection from system startup or background applications
- SYNOPSIS** /dev/msglog
- DESCRIPTION** Output from system startup (“rc”) scripts is directed to /dev/msglog, which dispatches it appropriately.
- ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr
Interface Stability	Stable

SEE ALSO `syslogd(1M)`, `syslog(3C)`, `attributes(5)`, `sysmsg(7D)`

NOTES In the current version of Solaris, /dev/msglog is an alias for /dev/sysmsg. In future versions of Solaris, writes to /dev/msglog may be directed into a more general logging mechanism such as `syslogd(1M)`.

`syslog(3C)` provides a more general logging mechanism than /dev/msglog and should be used in preference to /dev/msglog whenever possible.

NAME	msm – Microsoft Bus Mouse device interface
DESCRIPTION	The msm driver supports the Microsoft Bus Mouse. It allows applications to obtain information about the mouse's movements and the status of its buttons. The data is read in the Five Byte Packed Binary Format, also called MSC format.
FILES	/dev/msm
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `attributes(5)`

mt(7D)

NAME	mt – tape interface
DESCRIPTION	The files <code>rmt/*</code> refer to tape controllers and associated tape drives. The <code>labelit(1M)</code> command requires these magnetic tape file names to work correctly with the tape controllers. No other tape controller commands require these file names.
FILES	<code>/dev/rmt/*</code>
SEE ALSO	<code>labelit(1M)</code>

NAME	mtio – general magnetic tape interface
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ioctl.h> #include <sys/mtio.h></pre>
DESCRIPTION	<p>1/2", 1/4", 4mm, and 8mm magnetic tape drives all share the same general character device interface.</p> <p>There are two types of tape records: data records and end-of-file (EOF) records. EOF records are also known as tape marks and file marks. A record is separated by interrecord (or tape) gaps on a tape.</p> <p>End-of-recorded-media (EOM) is indicated by two EOF marks on 1/2" tape; by one EOF mark on 1/4", 4mm, and 8mm cartridge tapes.</p>
1/2" Reel Tape	<p>Data bytes are recorded in parallel onto the 9-track tape. Since it is a variable-length tape device, the number of bytes in a physical record may vary.</p> <p>The recording formats available (check specific tape drive) are 800 BPI, 1600 BPI, 6250 BPI, and data compression. Actual storage capacity is a function of the recording format and the length of the tape reel. For example, using a 2400 foot tape, 20 Mbyte can be stored using 800 BPI, 40 Mbyte using 1600 BPI, 140 Mbyte using 6250 BPI, or up to 700 Mbyte using data compression.</p>
1/4" Cartridge Tape	<p>Data is recorded serially onto 1/4" cartridge tape. The number of bytes per record is determined by the physical record size of the device. The I/O request size must be a multiple of the physical record size of the device. For QIC-11, QIC-24, and QIC-150 tape drives, the block size is 512 bytes.</p> <p>The records are recorded on tracks in a serpentine motion. As one track is completed, the drive switches to the next and begins writing in the opposite direction, eliminating the wasted motion of rewinding. Each file, including the last, ends with one file mark.</p> <p>Storage capacity is based on the number of tracks the drive is capable of recording. For example, 4-track drives can only record 20 Mbyte of data on a 450 foot tape; 9-track drives can record up to 45 Mbyte of data on a tape of the same length. QIC-11 is the only tape format available for 4-track tape drives. In contrast, 9-track tape drives can use either QIC-24 or QIC-11. Storage capacity is not appreciably affected by using either format. QIC-24 is preferable to QIC-11 because it records a reference signal to mark the position of the first track on the tape, and each block has a unique block number.</p> <p>The QIC-150 tape drives require DC-6150 (or equivalent) tape cartridges for writing. However, they can read other tape cartridges in QIC-11, QIC-24, or QIC-120 tape formats.</p>

mtio(7I)

8mm Cartridge Tape	Data is recorded serially onto 8mm helical scan cartridge tape. Since it is a variable-length tape device, the number of bytes in a physical record may vary. The recording formats available (check specific tape drive) are standard 2Gbyte, 5Gbyte, and compressed format.
4mm DAT Tape	Data is recorded either in Digital Data Storage (DDS) tape format or in Digital Data Storage, Data Compressed (DDS-DC) tape format. Since it is a variable-length tape device, the number of bytes in a physical record may vary. The recording formats available are standard 2Gbyte and compressed format.
Persistent Error Handling	<p>Persistent error handling is a modification of the current error handling behaviors, BSD and SVR4. With persistent error handling enabled, all tape operations after an error or exception will return immediately with an error. Persistent error handling can be most useful with asynchronous tape operations that use the <code>aioread(3AIO)</code> and <code>aiowrite(3AIO)</code> functions.</p> <p>To enable persistent error handling, the <code>ioctl MTIOCPERSISTENT</code> must be issued. If this <code>ioctl</code> succeeds, then persistent error handling is enabled and changes the current error behavior. This <code>ioctl</code> will fail if the device driver does not support persistent error handling.</p> <p>With persistent error handling enabled, all tape operations after an exception or error will return with the same error as the first command that failed; the operations will not be executed. An exception is some event that might stop normal tape operations, such as an End Of File (EOF) mark or an End Of Tape (EOT) mark. An example of an error is a media error. The <code>MTIOCLRERR</code> <code>ioctl</code> must be issued to allow normal tape operations to continue and to clear the error.</p> <p>Disabling persistent error handling returns the error behavior to normal SVR4 error handling, and will not occur until all outstanding operations are completed. Applications should wait for all outstanding operations to complete before disabling persistent error handling. Closing the device will also disable persistent error handling and clear any errors or exceptions.</p> <p>The <code>Read Operation</code> and <code>Write Operation</code> subsections contain more pertinent information regarding persistent error handling.</p>
Read Operation	The <code>read(2)</code> function reads the next record on the tape. The record size is passed back as the number of bytes read, provided it is not greater than the number requested. When a tape mark or end of data is read, a zero byte count is returned; all successive reads after the zero read will return an error and <code>errno</code> will be set to <code>EIO</code> . To move to the next file, an <code>MTFSF</code> <code>ioctl</code> can be issued before or after the read causing the error. This error handling behavior is different from the older BSD behavior, where another read will fetch the first record of the next tape file. If the BSD behavior is required, device names containing the letter <code>b</code> (for BSD behavior) in the final component should be used. If persistent error handling was enabled with either the BSD or SVR4 tape device behavior, all operations after this read error will return <code>EIO</code> errors until the <code>MTIOCLRERR</code> <code>ioctl</code> is issued. An <code>MTFSF</code> <code>ioctl</code> can then be issued.

Two successful successive reads that both return zero byte counts indicate EOM on the tape. No further reading should be performed past the EOM.

Fixed-length I/O tape devices require the number of bytes read to be a multiple of the physical record size. For example, 1/4" cartridge tape devices only read multiples of 512 bytes. If the blocking factor is greater than 64,512 bytes (minphys limit), fixed-length I/O tape devices read multiple records.

Most tape devices which support variable-length I/O operations may read a range of 1 to 65,535 bytes. If the record size exceeds 65,535 bytes, the driver reads multiple records to satisfy the request. These multiple records are limited to 65,534 bytes. Newer variable-length tape drivers may relax the above limitation and allow applications to read record sizes larger than 65,534. Refer to the specific tape driver man page for details.

Reading past logical EOT is transparent to the user. A read operation should never hit physical EOT.

Read requests that are lesser than a physical tape record are not allowed. Appropriate error is returned.

Write Operation

The `write(2)` function writes the next record on the tape. The record has the same length as the given buffer.

Writing is allowed on 1/4" tape at either the beginning of tape or after the last written file on the tape. With the Exabyte 8200, data may be appended only at the beginning of tape, before a filemark, or after the last written file on the tape.

Writing is not so restricted on 1/2", 4mm, and the other 8mm cartridge tape drives. Care should be used when appending files onto 1/2" reel tape devices, since an extra file mark is appended after the last file to mark the EOM. This extra file mark must be overwritten to prevent the creation of a null file. To facilitate write append operations, a space to the EOM `ioctl` is provided. Care should be taken when overwriting records; the erase head is just forward of the write head and any following records will also be erased.

Fixed-length I/O tape devices require the number of bytes written to be a multiple of the physical record size. For example, 1/4" cartridge tape devices only write multiples of 512 bytes.

Fixed-length I/O tape devices write multiple records if the blocking factor is greater than 64,512 bytes (minphys limit). These multiple writes are limited to 64,512 bytes. For example, if a write request is issued for 65,536 bytes using a 1/4" cartridge tape, two writes are issued; the first for 64,512 bytes and the second for 1024 bytes.

mtio(7I)

Most tape devices which support variable-length I/O operations may write a range of 1 to 65,535 bytes. If the record size exceeds 65,535 bytes, the driver writes multiple records to satisfy the request. These multiple records are limited to 65,534 bytes. As an example, if a write request for 65,540 bytes is issued, two records are written; one for 65,534 bytes followed by another record for 6 bytes. Newer variable-length tape drivers may relax the above limitation and allow applications to write record sizes larger than 65,534. Refer to the specific tape driver man page for details.

When logical EOT is encountered during a write, that write operation completes and the number of bytes successfully transferred is returned (note that a 'short write' may have occurred and not all the requested bytes would have been transferred. The actual amount of data written will depend on the type of device being used). The next write will return a zero byte count. A third write will successfully transfer some bytes (as indicated by the returned byte count, which again could be a short write); the fourth will transfer zero bytes, and so on, until the physical EOT is reached and all writes will fail with EIO.

When logical EOT is encountered with persistent error handling enabled, the current write may complete or be a short write. The next write will return a zero byte count. At this point an application should act appropriately for end of tape cleanup or issue yet another write, which will return the error ENOSPC. After clearing the exception with MTIOCLRERR, the next write will succeed (possibly short), followed by another zero byte write count, and then another ENOSPC error.

Allowing writes after LEOT has been encountered enables the flushing of buffers. However, it is strongly recommended to terminate the writing and close the file as soon as possible.

Seeks are ignored in tape I/O.

Close Operation

Magnetic tapes are rewound when closed, except when the "no-rewind" devices have been specified. The names of no-rewind device files use the letter n as the end of the final component. The no-rewind version of /dev/rmt/01 is /dev/rmt/01n. In case of error for a no-rewind device, the next open rewinds the device.

If the driver was opened for reading and a no-rewind device has been specified, the close advances the tape past the next filemark (unless the current file position is at EOM), leaving the tape correctly positioned to read the first record of the next file. However, if the tape is at the first record of a file it doesn't advance again to the first record of the next file. These semantics are different from the older BSD behavior. If BSD behavior is required where no implicit space operation is executed on close, the non-rewind device name containing the letter b (for BSD behavior) in the final component should be specified.

If data was written, a file mark is automatically written by the driver upon close. If the rewinding device was specified, the tape will be rewound after the file mark is written. If the user wrote a file mark prior to closing, then no file mark is written upon close. If a file positioning ioctl, like rewind, is issued after writing, a file mark is written before repositioning the tape.

All buffers are flushed on closing a tape device. Hence, it is strongly recommended that the application wait for all buffers to be flushed before closing the device. This can be done by writing a filemark via `MTWEOF`, even with a zero count.

Note that for 1/2" reel tape devices, two file marks are written to mark the EOM before rewinding or performing a file positioning ioctl. If the user wrote a file mark before closing a 1/2" reel tape device, the driver will always write a file mark before closing to insure that the end of recorded media is marked properly. If the non-rewinding device was specified, two file marks are written and the tape is left positioned between the two so that the second one is overwritten on a subsequent `open(2)` and `write(2)`.

If no data was written and the driver was opened for `WRITE-ONLY` access, one or two file marks are written, thus creating a null file.

After closing the device, persistent error handling will be disabled and any error or exception will be cleared.

IOCTLS Not all devices support all ioctls. The driver returns an `ENOTTY` error on unsupported ioctls.

The following structure definitions for magnetic tape ioctl commands are from `<sys/mtio.h>`.

The minor device byte structure is::

15	7	6	5	4	3	2	1	0
Unit #	BSD	Reserved	Density	Density	No rewind	Unit #		
Bits 7-15	behavior		Select	Select	on Close	Bits 0-1		

```

/*
 * Layout of minor device byte:
 */
#define MTUNIT(dev)      (((minor(dev) & 0xff80) >> 5) +
(minor(dev) & 0x3))
#define MT_NOREWIND     (1 <<2)
#define MT_DENSITY_MASK (3 <<3)
#define MT_DENSITY1     (0 <<3) /* Lowest density/format */
#define MT_DENSITY2     (1 <<3)
#define MT_DENSITY3     (2 <<3)
#define MT_DENSITY4     (3 <<3) /* Highest density/format */
#define MTMINOR(unit)   (((unit & 0x7fc) << 5) + (unit & 0x3))
#define MT_BSD          (1 <<6) /* BSD behavior on close */

/* Structure for MTIOCTOP - magnetic tape operation command */

```

mtio(7I)

```
struct mtop {
    short  mt_op;          /* operation */
    daddr_t mt_count;     /* number of operations */
};
```

The following operations of MTIOCTOP ioctl are supported:

MTWEOF	write an end-of-file record
MTFSF	forward space over file mark
MTBSF	backward space over file mark (1/2", 8mm only)
MTFSR	forward space to inter-record gap
MTBSR	backward space to inter-record gap
MTREW	rewind
MTOFFL	rewind and take the drive off-line
MTNOP	no operation, sets status only
MTRETEN	retension the tape (cartridge tape only)
MTERASE	erase the entire tape and rewind
MTEOM	position to EOM
MTNBSF	backward space file to beginning of file
MTSRSZ	set record size
MTGRSZ	get record size
MTLOAD	load the next tape cartridge into the tape drive

```
/* structure for MTIOCGET - magnetic tape get status command */

struct mtget {
    short  mt_type;       /* type of magtape device */
    /* the following two registers are device dependent */
    short  mt_dsreg;     /* "drive status" register */
    short  mt_erreg;     /* "error" register */
    /* optional error info. */
    daddr_t mt_resid;    /* residual count */
    daddr_t mt_fileno;   /* file number of current position */
    daddr_t mt_blkno;    /* block number of current position */
    ushort_t mt_flags;
    short  mt_bf;        /* optimum blocking factor */
};
/* structure for MTIOCGETDRIVETYPE - get tape config data command */
struct mtdrivetype_request {
    int  size;
    struct mtdrivetype *mtdtp;
};
struct mtdrivetype {
    char  name[64];      /* Name, for debug */
```

```

char    vid[25];                /* Vendor id and product id */
char    type;                  /* Drive type for driver */
int     bsize;                 /* Block size */
int     options;               /* Drive options */
int     max_rretries;          /* Max read retries */
int     max_wretries;          /* Max write retries */
uchar_t densities[MT_NDENSITIES]; /* density codes, low->hi */
uchar_t default_density;       /* Default density chosen */
uchar_t speeds[MT_NSPEEDS];    /* speed codes, low->hi */
ushort_t non_motion_timeout;   /* Seconds for non-motion */
ushort_t io_timeout;           /* Seconds for data to from tape */
ushort_t rewind_timeout;       /* Seconds to rewind */
ushort_t space_timeout;        /* Seconds to space anywhere */
ushort_t load_timeout;         /* Seconds to load tape and ready */
ushort_t unload_timeout;       /* Seconds to unload */
ushort_t erase_timeout;        /* Seconds to do long erase */
};

```

The `MTWEOF` ioctl is used for writing file marks to tape. Not only does this signify the end of a file, but also usually has the side effect of flushing all buffers in the tape drive to the tape medium. A zero count `MTWEOF` will just flush all the buffers and will not write any file marks. Because a successful completion of this tape operation will guarantee that all tape data has been written to the tape medium, it is recommended that this tape operation be issued before closing a tape device.

When spacing forward over a record (either data or EOF), the tape head is positioned in the tape gap between the record just skipped and the next record. When spacing forward over file marks (EOF records), the tape head is positioned in the tape gap between the next EOF record and the record that follows it.

When spacing backward over a record (either data or EOF), the tape head is positioned in the tape gap immediately preceding the tape record where the tape head is currently positioned. When spacing backward over file marks (EOF records), the tape head is positioned in the tape gap preceding the EOF. Thus the next read would fetch the EOF.

Record skipping does not go past a file mark; file skipping does not go past the EOM. After an `MTFSR <huge number>` command, the driver leaves the tape logically positioned *before* the EOF. A related feature is that EOFs remain pending until the tape is closed. For example, a program which first reads all the records of a file up to and including the EOF and then performs an `MTFSF` command will leave the tape positioned just after that same EOF, rather than skipping the next file.

The `MTNBSF` and `MTFSF` operations are inverses. Thus, an “`MTFSF -1`” is equivalent to an “`MTNBSF 1`”. An “`MTNBSF 0`” is the same as “`MTFSF 0`”; both position the tape device at the beginning of the current file.

`MTBSF` moves the tape backwards by file marks. The tape position will end on the beginning of the tape side of the desired file mark. An “`MTBSF 0`” will position the tape at the end of the current file, before the filemark.

mtio(7I)

MTBSR and MTFSR operations perform much like space file operations, except that they move by records instead of files. Variable-length I/O devices (1/2" reel, for example) space actual records; fixed-length I/O devices space physical records (blocks). 1/4" cartridge tape, for example, spaces 512 byte physical records. The status `ioctl` residual count contains the number of files or records not skipped.

MTOFFL rewinds and, if appropriate, takes the device off-line by unloading the tape. It is recommended that the device be closed after offlining and then re-opened after a tape has been inserted to facilitate portability to other platforms and other operating systems. Attempting to re-open the device with no tape will result in an error unless the `O_NDELAY` flag is used. (See `open(2)`.)

The `MTRETEN` retension `ioctl` applies only to 1/4" cartridge tape devices. It is used to restore tape tension, improving the tape's soft error rate after extensive start-stop operations or long-term storage.

MTERASE rewinds the tape, erases it completely, and returns to the beginning of tape. Erasing may take a long time depending on the device and/or tapes. For time details, refer to the the drive specific manual.

MTEOM positions the tape at a location just after the last file written on the tape. For 1/4" cartridge and 8mm tape, this is after the last file mark on the tape. For 1/2" reel tape, this is just after the first file mark but before the second (and last) file mark on the tape. Additional files can then be appended onto the tape from that point.

Note the difference between `MTBSF` (backspace over file mark) and `MTNBSF` (backspace file to beginning of file). The former moves the tape backward until it crosses an EOF mark, leaving the tape positioned *before* the file mark. The latter leaves the tape positioned *after* the file mark. Hence, "`MTNBSF n`" is equivalent to "`MTBSF (n+1)`" followed by "`MTFSF 1`". The 1/4" cartridge tape devices do not support `MTBSF`.

`MTRSRSZ` and `MTGRSZ` are used to set and get fixed record lengths. The `MTRSRSZ` `ioctl` allows variable length and fixed length tape drives that support multiple record sizes to set the record length. The `mt_count` field of the `mtop` struct is used to pass the record size to/from the `st` driver. A value of 0 indicates variable record size. The `MTRSRSZ` `ioctl` makes a variable-length tape device behave like a fixed-length tape device. Refer to the specific tape driver man page for details.

`MTLOAD` loads the next tape cartridge into the tape drive. This is generally only used with stacker and tower type tape drives which handle multiple tapes per tape drive. A tape device without a tape inserted can be opened with the `O_NDELAY` flag, in order to execute this operation.

The `MTIOCGET` get status `ioctl` call returns the drive ID (`mt_type`), sense key error (`mt_erreg`), file number (`mt_fileno`), optimum blocking factor (`mt_bf`) and record number (`mt_blkno`) of the last error. The residual count (`mt_resid`) is set to the number of bytes not transferred or files/records not spaced. The flags word (`mt_flags`) contains information such as whether the device is SCSI, whether it is a reel device, and whether the device supports absolute file positioning.

Persistent Error Handling IOCTLS and Asynchronous Tape Operations

The `MTIOCGETDRIVETYPE` get drivetype ioctl call returns the name of the tape drive as defined in `st.conf` (*name*), Vendor ID and model (*product*), ID (*vid*), type of tape device (*type*), block size (*bsize*), drive options (*options*), maximum read retry count (*max_rretries*), maximum write retry count (*max_wretries*), densities supported by the drive (*densities*), and default density of the tape drive (*default_density*).

<code>MTIOPERSISTENT</code>	enables/disables persistent error handling
<code>MTIOPERSISTENTSTATUS</code>	queries for persistent error handling
<code>MTIOCLRERR</code>	clears persistent error handling
<code>MTIOCGUARANTEEDORDER</code>	checks whether driver guarantees order of I/O's

The `MTIOPERSISTENT` ioctl enables or disables persistent error handling. It takes as an argument a pointer to an integer that turns it either on or off. If the ioctl succeeds, the desired operation was successful. It will wait for all outstanding I/O's to complete before changing the persistent error handling status. For example,

```
int on = 1;
ioctl(fd, MTIOPERSISTENT, &on);
int off = 0;
ioctl(fd, MTIOPERSISTENT, &off);
```

The `MTIOPERSISTENTSTATUS` ioctl enables or disables persistent error handling. It takes as an argument a pointer to an integer inserted by the driver. The integer can be either 1 if persistent error handling is 'on', or 0 if persistent error handling is 'off'. It will not wait for outstanding I/O's. For example,

```
int query;
ioctl(fd, MTIOPERSISTENTSTATUS, &query);
```

The `MTIOCLRERR` ioctl clears persistent error handling and allows tape operations to continual normally. This ioctl requires no argument and will always succeed, even if persistent error handling has not been enabled. It will wait for any outstanding I/O's before it clears the error.

The `MTIOCGUARANTEEDORDER` ioctl is used to determine whether the driver guarantees the order of I/O's. It takes no argument. If the ioctl succeeds, the driver will support guaranteed order. If the driver does not support guaranteed order, then it should not be used for asynchronous I/O with `libaio`. It will wait for any outstanding I/O's before it returns. For example,

```
ioctl(fd, MTIOCGUARANTEEDORDER)
```

See the Persistent Error Handling subsection above for more information on persistent error handling.

mtio(7I)

**Asynchronous and
State Change
IOCTLS**

MTIOCSTATE This ioctl blocks until the state of the drive, inserted or ejected, is changed. The argument is a pointer to a `mtio_state`, enum, whose possible enumerations are listed below. The initial value should be either the last reported state of the drive, or `MTIO_NONE`. Upon return, the enum pointed to by the argument is updated with the current state of the drive.

```
enum mtio_state {
MTIO_NONE      /* Return tape's current state */
MTIO_EJECTED   /* Tape state is "ejected" */
MTIO_INSERTED  /* Tape state is "inserted" */
;

```

When using asynchronous operations, most ioctls will wait for all outstanding commands to complete before they are executed.

**IOCTLS for
Multi-initiator
Configurations**

MTIOCRESERVE reserve the tape drive

MTIOCRELEASE revert back to the default behavior of reserve on open/release on close

MTIOCFORCERESERVE reserve the tape unit by breaking reservation held by another host

The **MTIOCRESERVE** ioctl reserves the tape drive such that it does not release the tape drive at close. This changes the default behavior of releasing the device upon close. Reserving the tape drive that is already reserved has no effect. For example,

```
ioctl(fd, MTIOCRESERVE);
```

The **MTIOCRELEASE** ioctl reverts back to the default behavior of reserve on open/release on close operation, and a release will occur during the next close. Releasing the tape drive that is already released has no effect. For example,

```
ioctl(fd, MTIOCRELEASE);
```

The **MTIOCFORCERESERVE** ioctl breaks a reservation held by another host, interrupting any I/O in progress by that other host, and then reserves the tape unit. This ioctl can be executed only with super-user privileges. It is recommended to open the tape device in `O_NDELAY` mode when this ioctl needs to be executed, otherwise the open will fail if another host indeed has it reserved. For example,

```
ioctl(fd, MTIOCFORCERESERVE);
```

**IOCTLS for
Handling Tape
Configuration
Options**

MTIOCSHORTFMK enables/disable support for writing short filemarks. This is specific to Exabyte drives.

MTIOCREADIGNOREILI enables/disable suppress incorrect length indicator support during reads

MTIOCREADIGNOREEOFs	enables/disable support for reading past two EOF marks which otherwise indicate End-Of-recording-Media (EOM) in the case of 1/2" reel tape drives
---------------------	---

The `MTIOCSHORTFMK` `ioctl` enables or disables support for short filemarks. This `ioctl` is only applicable to Exabyte drives which support short filemarks. As an argument, it takes a pointer to an integer. If 0 (zero) is the specified integer, then long filemarks will be written. If 1 is the specified integer, then short filemarks will be written. The specified tape behavior will be in effect until the device is closed.

For example:

```
int on = 1;
int off = 0;
/* enable short filemarks */
ioctl(fd, MTIOCSHORTFMK, &on);
/* disable short filemarks */
ioctl(fd, MTIOCSHORTFMK, &off);
```

Tape drives which do not support short filemarks will return an `errno` of `ENOTTY`.

The `MTIOCREADIGNOREILI` `ioctl` enables or disables the suppress incorrect length indicator (SILI) support during reads. As an argument, it takes a pointer to an integer. If 0 (zero) is the specified integer, SILI will not be used during reads and incorrect length indicator will not be suppressed. If 1 is the specified integer, SILI will be used during reads and incorrect length indicator will be suppressed. The specified tape behavior will be in effect until the device is closed.

For example:

```
int on = 1;
int off = 0;
ioctl(fd, MTIOCREADIGNOREILI, &on);
ioctl(fd, MTIOCREADIGNOREILI, &off);
```

The `MTIOCREADIGNOREEOFs` `ioctl` enables or disables support for reading past double EOF marks which otherwise indicate End-Of-recorded-media (EOM) in the case of 1/2" reel tape drives. As an argument, it takes a pointer to an integer. If 0 (zero) is the specified integer, then double EOF marks indicate End-Of-recorded-media (EOD). If 1 is the specified integer, the double EOF marks no longer indicate EOM, thus allowing applications to read past two EOF marks. In this case it is the responsibility of the application to detect end-of-recorded-media (EOM). The specified tape behavior will be in effect until the device is closed.

For example:

```
int on = 1;
int off = 0;
ioctl(fd, MTIOCREADIGNOREEOFs, &on);
```

mtio(7I)

```
ioctl(fd, MTIOREADIGNOREEOF, &off);
```

Tape drives other than 1/2" reel tapes will return an `errno` of `ENOTTY`.

EXAMPLES **EXAMPLE 1** Tape Positioning and Tape Drives

Suppose you have written three files to the non-rewinding 1/2" tape device, `/dev/rmt/0ln`, and that you want to go back and `dd(1M)` the second file off the tape. The commands to do this are:

```
mt -F /dev/rmt/0lbn bsf 3
mt -F /dev/rmt/0lbn fsf 1
dd if=/dev/rmt/0ln
```

To accomplish the same tape positioning in a C program, followed by a get status `ioctl`:

```
struct mtop mt_command;
struct mtget mt_status;
mt_command.mt_op = MTBSF;
mt_command.mt_count = 3;
ioctl(fd, MTIOCTOP, &mt_command);
mt_command.mt_op = MTFSF;
mt_command.mt_count = 1;
ioctl(fd, MTIOCTOP, &mt_command);
ioctl(fd, MTIOCGET, (char *)&mt_status);
```

or

```
mt_command.mt_op = MTNBSF;
mt_command.mt_count = 2;
ioctl(fd, MTIOCTOP, &mt_command);
ioctl(fd, MTIOCGET, (char *)&mt_status);
```

To get information about the tape drive:

```
struct mtdrivetype mtdt;
struct mtdrivetype_request mtreq;
mtreq.size = sizeof(struct mtdrivetype);
mtreq.mtdtp = &mtdt;
ioctl(fd, MTIOCGETDRIVETYPE, &mtreq);
```

FILES `/dev/rmt/<unit number><density>[<BSD behavior>][<no rewind>]`

Where *density* can be `l`, `m`, `h`, `u/c` (low, medium, high, ultra/compressed, respectively), the *BSD behavior* option is `b`, and the *no rewind* option is `n`.

For example, `/dev/rmt/0hbn` specifies unit 0, high density, BSD behavior and no rewind.

SEE ALSO `mt(1)`, `tar(1)`, `dd(1M)`, `open(2)`, `read(2)`, `write(2)`, `aioread(3AIO)`, `aiowrite(3AIO)`, `ar.h(3HEAD)`, `st(7D)`

1/4 Inch Tape Drive Tutorial

NAME	ncrs – SCSI host bus adapter driver
SYNOPSIS	<code>scsi@unit-address</code>
DESCRIPTION	<p>The <code>ncrs</code> host bus adapter driver is a SCSI-compliant nexus driver that supports the LSI Logic (formerly Symbios Logic or NCR) 53C810, 53C810A, 53C815, 53C820, 53C825, 53C825A, 53C860, 53C875, 53C875J, 53C876, and 53C895 SCSI (Small Computer Systems Interface) chips.</p> <p>The <code>ncrs</code> driver supports standard functions provided by the SCSI interface, including tagged and untagged queuing, Wide/Fast/Ultra/Ultra2 SCSI, and auto request sense. The <code>ncrs</code> driver does not support linked commands.</p>
PRECONFIGURATION	
Known Problems and Limitations	<ul style="list-style-type: none"> ■ The LSI BIOS and the Solaris <code>fdisk</code> program may be incompatible. To avoid problems, you should create an entry in the <code>FDISK</code> partition table using the DOS version of <code>FDISK</code> (or equivalent utility) before installing the Solaris software. To ensure your system will reboot following Solaris installation, create a DOS partition at least 1-cylinder in size that starts at cylinder 0. ■ Add-in cards containing 53C815, 53C820, 53C825, or 53C825A controllers must be used in bus-mastering PCI slots. PCI slots on dual PCI slot motherboards are generally bus-master capable. However, motherboards that contain three or more PCI slots, or motherboards that feature several embedded PCI controllers may contain PCI slots that are not bus-master capable. ■ PCI motherboards that feature LSI Logic SDMS BIOS and an embedded 53C810 or 53C810A controller may not be compatible with 53C82x add-in cards equipped with LSI Logic SDMS BIOS. To prevent conflicts, it may be necessary to upgrade the motherboard BIOS, the add-in card, or both. ■ Early PCI systems that are equipped with an 53C810 motherboard chip may contain unconnected interrupt pins. These systems cannot be used with Solaris software. ■ Wide-to-narrow target connections are not supported by Solaris software; as a result, you should not attempt to connect wide targets to narrow connectors on any of the supported devices. ■ If your adapter supports the LSI Logic SCSI configuration utility, the value of the host SCSI ID (found under the Adapter Setup menu) must be set to 7. (You can access the Symbios Logic SCSI configuration utility using Control-C.) ■ If you experience problems with old target devices, add the following to the <code>/kernel/drv/ncrs.conf</code> file: <pre>targetn-scsi-options = 0x0; where n is the ID of the failing target.</pre> ■ If you are using a Conner 1080S narrow SCSI drive, the system may display the following warnings:

ncrs(7D)

```
WARNING: /pci@0,0/pci1000,f@d (ncrs0):
invalid reselection (0,0)
WARNING: /pci@0,0/pci1000,f@d/sd@0,0 (sd0);
SCSI transport failed: 'reset: retrying command'
To suppress these warnings, disable tagged queuing in the ncrs.conf file.
```

- Pentium motherboards (Intel NX chipset) using P90 or slower processors may cause the ncrs driver to hang. If this occurs, the following messages are displayed on the console:

```
WARNING: /pci@0,0/pci1000,3@6 (ncrs0)
Unexpected DMA state:active dstat=c0<DMA-FIFO-empty,
master-data-parity-error>
This is an unrecoverable state and the system will not install using the ncrs driver.
```

- The ncrs driver supports the 53C875 chipset Revision 4, or later versions only. Pre-release versions of the chip are not supported.
- On rare occasions, use of an SDT7000/SDT9000 tape drive may result in the following message being displayed on the console:

```
Unexpected DMA state: ACTIVE. dstat=81<DMA-FIFO-empty,
illegal-instruction>
After the above message is displayed, the system and tape drive will recover and remain usable.
```

DRIVER CONFIGURATION

The ncrs host bus adapter driver is configured by defining the properties found in ncrs.conf. Properties in the ncrs.conf file that can be modified by the user include: scsi-options, target<n>-scsi-options, scsi-reset-delay, scsi-tag-age-limit, scsi-watchdog-tick, scsi-initiator-id, and ncrs-iomap. Properties in the ncrs.conf file override global SCSI settings.

The property target<n>-scsi-options overrides the scsi-options property value for target<n>, where <n> can vary from decimal 0 to 15. The ncrs driver supports the following SCSI options: SCSI_OPTIONS_DR (0x8), SCSI_OPTIONS_SYNC (0x20), SCSI_OPTIONS_TAG (0x80), SCSI_OPTIONS_FAST (0x100), SCSI_OPTIONS_WIDE (0x200), SCSI_OPTIONS_FAST20 (0x400), and SCSI_OPTIONS_FAST40 (0x800).

After periodic interval scsi-watchdog-tick, the ncrs driver searches through all current and disconnected commands for timeouts.

The scsi-tag-age-limit property represents the number of times that the ncrs driver attempts to allocate a tag ID that is currently in use after going through all tag IDs in a circular fashion. When encountering the same tag ID used scsi-tag-age-limit times, no additional commands are submitted to the target until all outstanding commands complete or timeout.

The ncrs-iomap property enables the driver to utilize IO mapping (rather than memory mapping) of registers.

Refer to scsi_hba_attach(9F) for details.

EXAMPLES | **EXAMPLE 1** A sample ncrs configuration file

Create a file called `/kernel/drv/ncrs.conf`, then add the following line:

```
scsi-options=0x78;
```

The above example disables tagged queuing, Fast/Ultra SCSI, and wide mode for all ncrs instances.

The following example disables an option for one specific ncrs device. See `driver.conf(4)` and `pci(4)` for more details.

```
name="ncrs" parent="/pci@1f,4000"
  unit-address="3"
  target1-scsi-options=0x58
  scsi-options=0x178 scsi-initiator-id=6;
```

In the example, the default initiator ID in OBP is 7; the change to ID 6 will occur at attach time. The `scsi-options` property is set for target 1 to 0x58 and all other targets set to 0x178. Note that it may be preferable to change the initiator ID in OBP.

The physical path name of the parent can be determined using the `/devices` tree or by following the link of the logical device name:

```
# ls -l /dev/rdisk/c0t0d0s0
lrwxrwxrwx 1 root  root    45 May 16 10:08 /dev/rdisk/c0t0d0s0 ->
  . . / . . /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

In the example above, the parent is `/pci@1f,4000` and the `unit-address` is the number bound to the `scsi@3` node.

To set `scsi-options` more specifically per target, do the following:

```
target1-scsi-options=0x78;
device-type-scsi-options-list =
"SEAGATE ST32550W", "seagate-scsi-options" ;
seagate-scsi-options = 0x58;
scsi-options=0x3f8;
```

With the exception of one specific disk type that has `scsi-options` set to 0x58, the example above sets `scsi-options` for target 1 to 0x78 and all other targets to 0x3f8.

The `scsi-options` properties that are specified per target ID have the highest precedence, followed by `scsi-options` per device type. Global `scsi-options` (for all ncrs instances) per bus have the lowest precedence.

To turn on IO mapping for all ncrs cards in the system, do the following:

```
ncrs-iomap=1;
```

The above action will noticeably slow the performance of the driver. You must reboot the system for the specified `scsi-options` to take effect.

ncrs(7D)

Driver Capabilities

To enable some driver features, the target driver must set capabilities in the `ncrs` driver. The following capabilities can be queried and modified by the target driver: `synchronous`, `tagged-qing`, `wide-xfer`, `auto-rqsense`, `qfull-retries`, and `qfull-retry-interval`. All other capabilities are query only.

The `tagged-qing`, `auto-rqsense`, `wide-xfer`, `disconnect`, and `Ultra/Ultra2` synchronous capabilities are enabled by default, and can be assigned binary (0 or 1) values only. The default value for `qfull-retries` is 10, while the default value for `qfull-retry-interval` is 100. The `qfull-retries` capability is a `uchar_t` (0 to 255), while `qfull-retry-interval` is a `ushort_t` (0 to 65535).

If a conflict exists between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only `whom != 0` is supported in the `scsi_ifsetcap(9F)` call. Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

The `ncrs` host bus adapter driver also supports hotplugging of targets using the `cfgadm` tool. Hotplug operations on the SCSI bus that hosts the root partition should not be performed. See the `cfgadm(1M)` man page for more information.

FILES

`/kernel/drv/ncrs` ELF kernel module
`/kernel/drv/ncrs.conf` Optional configuration file

ATTRIBUTES

See `attributes(5)` for descriptions of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Limited to PCI-based systems with Symbios 53C810, 53C810A, 53C815, 53C820, 53C825, 53C825A, 53C860, 53C875, 53C875J, 53C876, and 53C895 SCSI I/O processors.

SEE ALSO

`prtconf(1M)`, `driver.conf(4)`, `pci(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

ANSI Small Computer System Interface-2 (SCSI-2)

Symbios Logic Inc., *SYM53C895 PCI-Ultra2 SCSI I/O Processor With LVDlink*

Symbios Logic Inc., *SYM53C875 PCI-SCSI I/O Processor With Fast-20*

Symbios Logic Inc., *SYM53C825A PCI-SCSI I/O Processor*

Symbios Logic Inc., *SYM53C810A PCI-SCSI I/O Processor*

DIAGNOSTICS

The messages described below are logged and may also appear on the system console.

Device is using a hilevel intr

The device was configured with an interrupt level that cannot be used with this ncrs driver. Check the PCI device.

map setup failed

The driver was unable to map device registers; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

glm_script_alloc failed

The driver was unable to load the SCRIPTS for the SCSI processor; check for bad hardware. Driver did not attach to device; SCSI devices will be inaccessible.

cannot map configuration space

The driver was unable to map in the configuration registers. Check for bad hardware. SCSI devices will be inaccessible

attach failed

The driver was unable to attach; usually preceded by another warning that indicates why attach failed. These can be considered hardware failures.

SCSI bus DATA IN phase parity error

The driver detected parity errors on the SCSI bus.

SCSI bus MESSAGE IN phase parity error

The driver detected parity errors on the SCSI bus.

SCSI bus STATUS phase parity error

The driver detected parity errors on the SCSI bus.

Unexpected bus free

Target disconnected from the bus without notice. Check for bad hardware.

Disconnected command timeout for Target <id>.<lun>

A timeout occurred while target *id/lun* was disconnected. This is usually a target firmware problem. For tagged queuing targets, <*n*> commands were outstanding when the timeout was detected.

Disconnected tagged cmd(s) (<*n*>) timeout for Target <id>.<lun>

A timeout occurred while target *id/lun* was disconnected. This is usually a target firmware problem. For tagged queuing targets, <*n*> commands were outstanding when the timeout was detected.

Connected command timeout for Target <id>.<lun>

This is usually a SCSI bus problem. Check cables and termination.

ncrs(7D)

```
Target <id> reducing sync. transfer rate
```

A data transfer hang or DATA-IN phase parity error was detected. The driver attempts to eliminate this problem by reducing the data transfer rate.

```
Target <id> reverting to async. mode
```

A second data transfer hang was detected for this target. The driver attempts to eliminate this problem by reducing the data transfer rate.

```
Target <id> disabled wide SCSI mode
```

A second data phase hang was detected for this target. The driver attempts to eliminate this problem by disabling wide SCSI mode.

```
auto request sense failed
```

An attempt to start an auto request packet failed. Another auto request packet may already be in transport.

```
invalid reselection (<id>.<lun>)
```

A reselection failed; target accepted abort or reset, but still tries to reconnect. Check for bad hardware.

```
invalid intcode
```

The SCRIPTS processor generated an invalid SCRIPTS interrupt. Check for bad hardware.

NOTES

The ncrs hardware (53C875) supports Wide, Fast, and Ultra SCSI mode. The maximum SCSI bandwidth is 40 MB/sec.

The ncrs hardware (53C895) supports Wide, Fast, Ultra and Ultra2 SCSI mode using a LVD bus. The maximum SCSI bandwidth is 80 MB/second.

The ncrs driver exports properties indicating the negotiated transfer speed per target (`target<n>-sync-speed`), whether wide bus is supported (`target<n>-wide`) for that particular target (`target<n>-scsi-options`), and whether tagged queuing has been enabled (`target<n>-TQ`). The `sync-speed` property value indicates the data transfer rate in KB/sec. The `target<n>-TQ` and the `target<n>-wide` property have value 1 (to indicate that the corresponding capability is enabled for that target), or 0 (to indicate that the capability is disabled for that target). See `prtconf(1M)` (verbose option) for details on viewing the ncrs properties.

```
scsi, instance #0
  Driver properties:
    name <target6-TQ> length <4>
      value <0x00000000>.
    name <target6-wide> length <4>
      value <0x00000000>.
    name <target6-sync-speed> length <4>
      value <0x00002710>.
    name <target1-TQ> length <4>
```

```
    value <0x00000001>.  
name <target1-wide> length <4>  
    value <0x00000000>.  
name <target1-sync-speed> length <4>  
    value <0x00002710>.  
name <target0-TQ> length <4>  
    value <0x00000001>.  
name <target0-wide> length <4>  
    value <0x00000001>.  
name <target0-sync-speed> length <4>  
    value <0x00009c40>.  
name <scsi-options> length <4>  
    value <0x000007f8>.  
name <scsi-watchdog-tick> length <4>  
    value <0x0000000a>.  
name <scsi-tag-age-limit> length <4>  
    value <0x00000002>.  
name <scsi-reset-delay> length <4>  
    value <0x00000bb8>.  
name <latency-timer> length <4>  
    value <0x00000088>.  
name <cache-line-size> length <4>  
    value <0x00000010>.
```

null(7D)

NAME	null – the null file, also called the null device
SYNOPSIS	<code>/dev/null</code>
DESCRIPTION	<p>Data written on the null special file, <code>/dev/null</code>, is discarded.</p> <p>Reads from a null special file always return 0 bytes.</p> <p>Mapping a null special file creates an address reservation of a length equal to the length of the mapping, and rounded up to the nearest page size as returned by <code>sysconf(3C)</code>. No resources are consumed by the reservation. Mappings can be placed in the resulting address range via subsequent calls to <code>mmap</code> with the <code>-MAP_FIXED</code> option set.</p>
FILES	<code>/dev/null</code>
SEE ALSO	<code>mmap(2)</code> , <code>sysconf(3C)</code>

NAME	objfs – Kernel object filesystem				
DESCRIPTION	<p>The <code>objfs</code> filesystem describes the state of all modules currently loaded by the kernel. It is mounted during boot at <code>/system/object</code>.</p> <p>The contents of the filesystem are dynamic and reflect the current state of the system. Each module is represented by a directory containing a single file, 'object.' The object file is a read only ELF file which contains information about the object loaded in the kernel.</p> <p>The kernel may load and unload modules dynamically as the system runs. As a result, applications may observe different directory contents in <code>/system/object</code> if they repeatedly rescan the directory. If a module is unloaded, its associated <code>/system/object</code> files disappear from the hierarchy and subsequent attempts to open them, or to read files opened before the module unloaded, elicits an error.</p>				
FILES	<p><code>/system/object</code> Mount point for objfs file system</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Private</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Private
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Private				
SEE ALSO	<code>vfstab(4)</code>				
NOTES	The content of the ELF files is private to the implementation and subject to change without notice.				

ocf_ibutton(7D)

NAME	ocf_ibutton – iButton Smart Card terminal driver
DESCRIPTION	The <code>ocf_ibutton</code> smart card terminal driver is an OpenCard Framework (OCF)-compliant terminal driver for the Dallas Semiconductor iButton reader.
APPLICATION PROGRAMMING INTERFACE	The <code>ocf_ibutton</code> smart card terminal driver is part of the OCF framework stack and is started by the OCF startup script. The iButton reader requires a host serial port and is accessed through the character-special devices.
FILES	<code>/usr/share/lib/smartcard/ibutton.jar</code> Java-archived driver class files <code>/dev/cua/x</code> Asynchronous serial line using port x
SEE ALSO	<code>ports(1M)</code> , <code>smartcard(1M)</code> , <code>smartcard(5)</code>

NAME	ohci – OpenHCI host controller driver						
SYNOPSIS	usb@unit-address						
DESCRIPTION	<p>The ohci driver is a USBA (Solaris USB Architecture) compliant nexus driver that supports the <i>Open Host Controller Interface Specification 1.1</i>, an industry standard developed by Compaq, Microsoft, and National Semiconductor.</p> <p>The ohci driver supports bulk, interrupt, control and isochronous transfers.</p>						
FILES	<p>/kernel/drv/ohci 32-bit x86 ELF kernel module</p> <p>/kernel/drv/amd64/ohci 64-bit x86 ELF kernel module</p> <p>/kernel/drv/sparcv9/ohci 64-bit SPARC ELF kernel module</p> <p>/kernel/drv/ohci.conf driver configuration file</p>						
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>SPARC, x86, PCI-based systems</td> </tr> <tr> <td>Availability</td> <td>SUNWusb</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC, x86, PCI-based systems	Availability	SUNWusb
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	SPARC, x86, PCI-based systems						
Availability	SUNWusb						
SEE ALSO	<p>attributes(5), ehci(7D), hubd(7D), uhci(7D), usba(7D)</p> <p><i>Writing Device Drivers</i></p> <p><i>Universal Serial Bus Specification 2.0</i></p> <p><i>Open Host Controller Interface Specification for USB 1.0a</i></p> <p><i>System Administration Guide: Basic Administration</i></p> <p>http://www.sun.com/io</p>						
DIAGNOSTICS	<p>All host controller errors are passed to the client drivers. Root hub errors are documented in hubd(7D).</p> <p>In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:</p> <p>WARNING: <device path> <ohci><instance number>>: Error message...</p> <p>Unrecoverable USB Hardware Error. There was an unrecoverable USB hardware error reported by the OHCI Controller. Please reboot the system. If this problem persists, contact your system vendor.</p>						

ohci(7D)

No SOF interrupts have been received. This OHCI USB controller is unusable.
The USB hardware is not generating Start Of Frame interrupts. Please reboot the system. If this problem persists, contact your system vendor.

NAME	openprom – PROM monitor configuration interface						
SYNOPSIS	<pre>#include <sys/fcntl.h> #include <sys/types.h> #include <sys/openpromio.h> open("/dev/openprom", mode);</pre>						
DESCRIPTION	<p>The internal encoding of the configuration information stored in EEPROM or NVRAM varies from model to model, and on some systems the encoding is “hidden” by the firmware. The openprom driver provides a consistent interface that allows a user or program to inspect and modify that configuration, using <code>ioctl(2)</code> requests. These requests are defined in <code><sys/openpromio.h></code>:</p> <pre>struct openpromio { uint_t oprom_size; /* real size of following data */ union { char b[1]; /* NB: Adjacent, Null terminated */ int i; } opio_u; }; #define oprom_array opio_u.b /* property name/value array */ #define oprom_node opio_u.i /* nodeid from navigation config-ops */ #define oprom_len opio_u.i /* property len from OPROMGETPROPLEN */ #define OPROMMAXPARAM 32768 /* max size of array (advisory) */</pre> <p>For all <code>ioctl(2)</code> requests, the third parameter is a pointer to a <code>struct openpromio</code>. All property names and values are null-terminated strings; the value of a numeric option is its ASCII representation.</p> <p>For the raw <code>ioctl(2)</code> operations shown below that explicitly or implicitly specify a nodeid, an error may be returned. This is due to the removal of the node from the firmware device tree by a Dynamic Reconfiguration operation. Programs should decide if the appropriate response is to restart the scanning operation from the beginning or terminate, informing the user that the tree has changed.</p>						
IOCTLS	<table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">OPROMGETOPT</td> <td>This <code>ioctl</code> takes the null-terminated name of a property in the <code>oprom_array</code> and returns its null-terminated value (overlying its name). <code>oprom_size</code> should be set to the size of <code>oprom_array</code>; on return it will contain the size of the returned value. If the named property does not exist, or if there is not enough space to hold its value, then <code>oprom_size</code> will be set to zero. See BUGS below.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">OPROMSETOPT</td> <td>This <code>ioctl</code> takes two adjacent strings in <code>oprom_array</code>; the null-terminated property name followed by the null-terminated value.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">OPROMSETOPT2</td> <td>This <code>ioctl</code> is similar to <code>OPROMSETOPT</code>, except that it uses the difference between the actual user array size and the length of the property name plus its null terminator.</td> </tr> </table>	OPROMGETOPT	This <code>ioctl</code> takes the null-terminated name of a property in the <code>oprom_array</code> and returns its null-terminated value (overlying its name). <code>oprom_size</code> should be set to the size of <code>oprom_array</code> ; on return it will contain the size of the returned value. If the named property does not exist, or if there is not enough space to hold its value, then <code>oprom_size</code> will be set to zero. See BUGS below.	OPROMSETOPT	This <code>ioctl</code> takes two adjacent strings in <code>oprom_array</code> ; the null-terminated property name followed by the null-terminated value.	OPROMSETOPT2	This <code>ioctl</code> is similar to <code>OPROMSETOPT</code> , except that it uses the difference between the actual user array size and the length of the property name plus its null terminator.
OPROMGETOPT	This <code>ioctl</code> takes the null-terminated name of a property in the <code>oprom_array</code> and returns its null-terminated value (overlying its name). <code>oprom_size</code> should be set to the size of <code>oprom_array</code> ; on return it will contain the size of the returned value. If the named property does not exist, or if there is not enough space to hold its value, then <code>oprom_size</code> will be set to zero. See BUGS below.						
OPROMSETOPT	This <code>ioctl</code> takes two adjacent strings in <code>oprom_array</code> ; the null-terminated property name followed by the null-terminated value.						
OPROMSETOPT2	This <code>ioctl</code> is similar to <code>OPROMSETOPT</code> , except that it uses the difference between the actual user array size and the length of the property name plus its null terminator.						

openprom(7D)

	OPROMNXTOPT	This ioctl is used to retrieve properties sequentially. The null-terminated name of a property is placed into <i>oprom_array</i> and on return it is replaced with the null-terminated name of the next property in the sequence, with <i>oprom_size</i> set to its length. A null string on input means return the name of the first property; an <i>oprom_size</i> of zero on output means there are no more properties.
	OPROMNXT OPROMCHILD OPROMGETPROP OPROMNXTPROP	These ioctls provide an interface to the raw <i>config_ops</i> operations in the PROM monitor. One can use them to traverse the system device tree; see <code>prtconf(1M)</code> .
	OPROMGETPROPLEN	This ioctl provides an interface to the <i>property length</i> raw config op. It takes the name of a property in the buffer, and returns an integer in the buffer. It returns the integer -1 if the property does not exist; 0 if the property exists, but has no value (a boolean property); or a positive integer which is the length of the property as reported by the PROM monitor. See BUGS below.
	OPROMGETVERSION	This ioctl returns an arbitrary and platform-dependent NULL-terminated string in <i>oprom_array</i> , representing the underlying version of the firmware.
ERRORS	EAGAIN	There are too many opens of the <code>/dev/openprom</code> device.
	EFAULT	A bad address has been passed to an <code>ioctl(2)</code> routine.
	EINVAL	The size value was invalid, or (for <code>OPROMSETOPT</code>) the property does not exist, or an invalid ioctl is being issued, or the ioctl is not supported by the firmware, or the <code>nodeid</code> specified does not exist in the firmware device tree.
	ENOMEM	The kernel could not allocate space to copy the user's structure.
	EPERM	Attempts have been made to write to a read-only entity, or read from a write only entity.
	ENXIO	Attempting to open a non-existent device.
EXAMPLES	EXAMPLE 1 <i>oprom_array</i> Data Allocation and Reuse	
		The following example shows how the <i>oprom_array</i> is allocated and reused for data returned by the driver.
		<pre>/* * This program opens the openprom device and prints the platform * name (root node name property) and the prom version. */</pre>

EXAMPLE 1 *oprom_array* Data Allocation and Reuse (Continued)

```

* NOTE: /dev/openprom is readable only by user 'root' or group 'sys'.
*/
#include <stdio.h>
#include <string.h>
#include <fcntl.h>
#include <errno.h>
#include <unistd.h>
#include <stdlib.h>
#include <sys/openpromio.h>
#define min(a, b) (a < b ? a : b)
#define max(a, b) (a > b ? a : b)
#define MAXNAMESZ 32 /* Maximum property *name* size */
#define BUFSZ 1024 /* A Handy default buffer size */
#define MAXVALSZ (BUFSZ - sizeof (int))
static char *promdev = "/dev/openprom";
/*
* Allocate an openpromio structure big enough to contain
* a bufsize'd oprom_array. Zero out the structure and
* set the oprom_size field to bufsize.
*/
static struct openpromio *
opp_zalloc(size_t bufsize)
{
    struct openpromio *opp;
    opp = malloc(sizeof (struct openpromio) + bufsize);
    (void) memset(opp, 0, sizeof (struct openpromio) + bufsize);
    opp->oprom_size = bufsize;
    return (opp);
}
/*
* Free a 'struct openpromio' allocated by opp_zalloc
*/
static void
opp_free(struct openpromio *opp)
{
    free(opp);
}
/*
* Get the peer node of the given node. The root node is the peer of zero.
* After changing nodes, property lookups apply to that node. The driver
* 'remembers' what node you are in.
*/
static int
peer(int nodeid, int fd)
{
    struct openpromio *opp;
    int i;
    opp = opp_zalloc(sizeof (int));
    opp->oprom_node = nodeid;
    if (ioctl(fd, OPROMNEXT, opp) < 0) {
        perror("OPROMNEXT");
        exit(1);
    }
    i = opp->oprom_node;

```

EXAMPLE 1 *oprom_array* Data Allocation and Reuse (Continued)

```

    opp_free(opp);
    return(i);
}
int
main(void)
{
    struct openpromio *opp;
    int fd, proplen;
    size_t buflen;
    if ((fd = open(promdev, O_RDONLY)) < 0) {
        fprintf(stderr, "Cannot open openprom device\n");
        exit(1);
    }
    /*
     * Get and print the length and value of the
     * root node 'name' property
     */
    (void) peer(0, fd);      /* Navigate to the root node */
    /*
     * Allocate an openpromio structure sized big enough to
     * take the string "name" as input and return the int-sized
     * length of the 'name' property.
     * Then, get the length of the 'name' property.
     */
    buflen = max(sizeof (int), strlen("name") + 1);
    opp = opp_zalloc(buflen);
    (void) strcpy(opp->oprom_array, "name");
    if (ioctl(fd, OPROMGETPROPLEN, opp) < 0) {
        perror("OPROMGETPROPLEN");
        /* exit(1); */
        proplen = 0;      /* down-rev driver? */
    } else
        proplen = opp->oprom_len;
    opp_free(opp);
    if (proplen == -1) {
        printf("'name' property does not exist!\n");
        exit (1);
    }
    /*
     * Allocate an openpromio structure sized big enough
     * to take the string 'name' as input and to return
     * 'proplen + 1' bytes. Then, get the value of the
     * 'name' property. Note how we make sure to size the
     * array at least one byte more than the returned length
     * to guarantee NULL termination.
     */
    buflen = (proplen ? proplen + 1 : MAXVALSZ);
    buflen = max(buflen, strlen("name") + 1);
    opp = opp_zalloc(buflen);
    (void) strcpy(opp->oprom_array, "name");
    if (ioctl(fd, OPROMGETPROP, opp) < 0) {
        perror("OPROMGETPROP");
        exit(1);
    }
}

```


EXAMPLE 1 *oprom_array* Data Allocation and Reuse (Continued)

```

if (opp->oprom_size != 0)
    printf("Platform name <%s> property len <%d>\n",
        opp->oprom_array, proplen);
opp_free(opp);
/*
 * Allocate an openpromio structure assumed to be
 * big enough to get the 'prom version string'.
 * Get and print the prom version.
 */
opp_zalloc(MAXVALSZ);
opp->oprom_size = MAXVALSZ;
if (ioctl(fd, OPROMGETVERSION, opp) < 0) {
    perror("OPROMGETVERSION");
    exit(1);
}
printf("Prom version <%s>\n", opp->oprom_array);
opp_free(opp);
(void) close(fd);
return (0);
}

```

FILES /dev/openprom PROM monitor configuration interface

SEE ALSO eeprom(1M), monitor(1M), prtconf(1M), ioctl(2), mem(7D)

BUGS There should be separate return values for non-existent properties as opposed to not enough space for the value.

An attempt to set a property to an illegal value results in the PROM setting it to some legal value, with no error being returned. An OPROMGETOPT should be performed after an OPROMSETOPT to verify that the set worked.

Some PROMS *lie* about the property length of some string properties, omitting the NULL terminator from the property length. The openprom driver attempts to *transparently* compensate for these bugs when returning property values by NULL terminating an extra character in the user buffer if space is available in the user buffer. This extra character is excluded from the *oprom_size* field returned from OPROMGETPROP and OPROMGETOPT and excluded in the *oprom_len* field returned from OPROMGETPROPLEN but is returned in the user buffer from the calls that return data, if the user buffer is allocated at least one byte larger than the property length.

pcata(7D)

NAME	pcata – PCMCIA ATA card device driver									
SYNOPSIS	<pre>pcata@socket#:a -u pcata@socket#:a -u,raw</pre>									
DESCRIPTION	<p>The PCMCIA ATA card device driver supports PCMCIA ATA disk and flash cards that follow the following standards:</p> <ul style="list-style-type: none">■ PC card 2.01 compliance (MBR+fdisk table required for all platforms).■ PC card ATA 2.01 compliance.■ PC card services 2.1 compliance. <p>The driver supports standard PCMCIA ATA cards that contain a Card Information Structure (CIS). For PCMCIA, nodes are created in <code>/devices</code> that include the socket number as one component of the device name referred to by the node. However, the names in <code>/dev</code>, <code>/dev/dsk</code>, and <code>/dev/rdisk</code> follow the current conventions for ATA devices, which do not encode the socket number in any part of the name. For example, you may have the following:</p> <table border="1"><thead><tr><th>Platform</th><th>/devices name</th><th>/dev/dsk name</th></tr></thead><tbody><tr><td>x86</td><td><code>/devices/isa/pcic@1,3e0</code> <code>/disk@0:a</code></td><td><code>/dev/dsk/c1d0s0</code></td></tr><tr><td>SPARC</td><td><code>/devices/iommu@f,e0000000</code> <code>/sbus@f,e0001000</code> <code>/SUNW,pcmcia@3,0</code> <code>/disk@0:a</code></td><td><code>/dev/dsk/c1d0s0</code></td></tr></tbody></table>	Platform	/devices name	/dev/dsk name	x86	<code>/devices/isa/pcic@1,3e0</code> <code>/disk@0:a</code>	<code>/dev/dsk/c1d0s0</code>	SPARC	<code>/devices/iommu@f,e0000000</code> <code>/sbus@f,e0001000</code> <code>/SUNW,pcmcia@3,0</code> <code>/disk@0:a</code>	<code>/dev/dsk/c1d0s0</code>
Platform	/devices name	/dev/dsk name								
x86	<code>/devices/isa/pcic@1,3e0</code> <code>/disk@0:a</code>	<code>/dev/dsk/c1d0s0</code>								
SPARC	<code>/devices/iommu@f,e0000000</code> <code>/sbus@f,e0001000</code> <code>/SUNW,pcmcia@3,0</code> <code>/disk@0:a</code>	<code>/dev/dsk/c1d0s0</code>								
PRECONFIGURATION	<p>When a PC Card ATA device is recognized, the <code>pcata</code> driver is automatically loaded, IRQs allocated, devices nodes created, and special files created (if they do not already exist).</p>									
Known Problems and Limitations	<ul style="list-style-type: none">■ <code>vold</code> does not support <code>pcata</code>. File systems must be mounted manually.■ You need to unmount the file system before removing the disk.■ The <code>ufs</code> file systems on removable media (PC Card ATA) should have one of the <code>onerror={panic, lock, umount}</code> mount options set.									
CONFIGURATION	<p>Configuration topics include initial installation and configuration, identifying an unrecognized device, special files and hot-plugging.</p>									
Initial Installation and Configuration	<ol style="list-style-type: none">1. Install the Solaris software.2. Boot the system.3. Insert the PC card ATA device.									
Identifying an Unrecognized Device	<p>If you insert a PC card ATA device and it is not recognized (no special files created), use the <code>prtconf</code> command to identify the problem.</p> <ol style="list-style-type: none">1. Run the <code>prtconf -D</code> command to see if your <code>pcata</code> card is recognized. (A recognized device will appear at the end of the <code>prtconf</code> output. For example:									

```
# prtconf -D
. . .
pcic, instance #0 (driver name: pcic)
. . .
disk, instance #0
```

2. If pcata does not appear in the prtconf output, there is a problem with the PC card adapter configuration or with the hardware. Check to see whether the problem is with the card or the adapter by trying to use the card on another machine and by seeing if it works on the same machine using DOS.

Special Files

For PC card devices, nodes are created in `/devices` that include the socket number as one component of a device name that the node refers to. However, the `/prtcd/dev` names and the names in `/dev/dsk` and `/dev/rdisk` do follow the current convention for ATA devices, which do not encode the socket number in any part of the name.

Hot-Plugging

- If you want to remove the disk, you must unmount the file system.
- Use the `mkfs_pcfs(1M)` command to create a `pcfs` file system:

```
# mkfs -F pcfs /dev/rdisk/c#d#p0:d
```

- To mount a `pcfs` file system, type:

```
# mount -F pcfs /dev/dsk/c#d#p0:c /mnt
```

- If you want to create a `ufs` file system, use the `newfs` command and type:

```
# newfs /dev/rdisk/c#d#s#
```

- To mount a `ufs` file system, type:

```
# mount -F ufs /dev/dsk/c#d#s# /mnt
```

- To create a Solaris partition, run the `format` command and go to the Partition menu. For more information, see the `format(1M)` man page.

FILES

`/kernel/drv/pcata` pcata driver

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWpsdpr

SEE ALSO

`format(1M)`, `mount(1M)`, `newfs(1M)`, `pcmcia(4)`, `attributes(5)`, `pcfs(7FS)`

pcelx(7D)

NAME	pcelx – 3COM EtherLink III PCMCIA Ethernet Adapter
SYNOPSIS	network@<socket>:pcelx<socket>
DESCRIPTION	<p>The pcelx driver supports the 3COM EtherLink III PCMCIA PC Card as a standard Ethernet type of device conforming to the DLPI interface specification. The driver supports the <i>hot-plugging</i> of the PC Card.</p> <p>The PPA (Physical Point of Attachment) is defined by the socket number the PC Card is inserted in. This means that for IP use, the PC Card should always be plugged into the same socket that the network interface was initially brought up on or else a network reconfiguration should be done to take down the old interface and bring up the new one.</p> <p>The 3C589, 3C589B, and 3C589C versions of the PC Card are supported on the x86 platform. The 3C589B and 3C589C are supported on the SPARC platform.</p>
PRECONFIGURE	<ul style="list-style-type: none">■ For IBM ThinkPad 760E series systems and systems using the TI PCI1130 PCI-to-CardBus chip (such as the Dell Latitude XPi CD) only: Before bringing the system onto the network, put the PC Card into 8-bit mode by creating a file called /kernel/drv/pcelx.conf containing force-8bit=1;.■ It is not possible to boot or install the Solaris software using a 3Com EtherLink III PC Card device.■ If the 3Com PC Card device is recognized, the pcelx driver is automatically loaded, ports and IRQs allocated, and special files created (if they don't already exist). No manual configuration of the hardware is necessary or possible.
Known Problems and Limitations	Network services are automatically started when the system is booted. These services are not started when a network interface is added or shut down after the system has been brought up.
CONFIGURATION	Configuration procedures include initial installation and configuration, identifying an unrecognized card and configuring two or more cards.
Initial Installation and Configuration	<p>For initial installation and configuration, do the following steps:</p> <ol style="list-style-type: none">1. Install the Solaris software.2. Boot the system.3. Insert the 3Com EtherLink III PC Card device.
Identifying an Unrecognized Card	<p>If you insert a 3C589 card and it is not recognized (no special files created), use the prtconf command and attempt to identify the problem:</p> <ol style="list-style-type: none">1. Become root.2. Run the prtconf -D command to see if your 3C589 card is recognized. A recognized device will appear in the prtconf output. For example: <pre># prtconf -D . . . pcic, instance #0 (driver name: pcic) . . .</pre>

```
network, instance #0 (driver name: pcelx)
```

3. If `pcelx` does not appear in the `prtconf` output, there is a problem with the PC Card adapter configuration or with the hardware. You can determine whether the problem is with the card or the adapter by attempting to use the card on another machine or by using it on the same machine using DOS.

Configuring Two or More Cards

Because the 3C589 card is not supported during Solaris installation, you must update network configuration files before it can be used as a network interface:

1. Create a `/etc/hostname.pcelx#` file (where # is a socket number) to specify the host name to be associated with this interface.
2. Add an IP address for the new host name to the `/etc/inet/hosts` file.
3. Ensure that the associated network is listed in `/etc/inet/netmasks`.
4. Ensure that the Name Service Switch `/etc/nsswitch.conf` configuration file includes the network and local services you need.
5. Reboot the system.

Special Files

Device naming in `/dev` follows standard LAN device naming with the exception that the PPA (physical point of attachment) unit number is the socket where the card resides, not the instance. For the `pcelx` driver, `/dev/pcelx0` (or PPA 0 of `/dev/pcelx`) is the card in socket 0, while a card in socket 1 is `/dev/pcelx1` (or PPA 1 of `/dev/pcelx`).

Hot Plugging

If you remove the 3C589 card, any information you send is discarded, and no error messages are given.

When you reinsert the card in the same socket, the device operates normally. The behavior is similar to temporarily disconnecting the device from the network.

FILES	<code>/kernel/drv/pcelx</code>	<code>pcelx</code> driver
	<code>/dev/pcelx</code>	DLPI Style 2 device
	<code>/dev/pcelxn</code>	DLPI Style 1 device where: <i>n</i> is the PCMCIA physical socket number.

SEE ALSO

`pcmcia(4)`

pcfs(7FS)

NAME	pcfs – DOS formatted file system
SYNOPSIS	<pre>#include <sys/param.h> #include <sys/mount.h> #include <sys/fs/pc_fs.h> int mount(const char *spec, const char *dir, int mflag, "pcfs", struct pcfs_args, struct *pc_argp, sizeof (struct pcfs_args));</pre>
DESCRIPTION	<p>pcfs is a file system type that enables direct access to files on DOS formatted disks from within the SunOS operating system.</p> <p>Once mounted, pcfs provides standard SunOS file operations and semantics. Using pcfs, you can create, delete, read, and write files on a DOS formatted disk. You can also create and delete directories and list files in a directory.</p> <p>pcfs supports FAT12 (floppies) and FAT16 and FAT32 file systems.</p> <p>The pcfs file system contained on the block special file identified by <i>spec</i> is mounted on the directory identified by <i>dir</i>. <i>spec</i> and <i>dir</i> are pointers to pathnames. <i>mflag</i> specifies the mount options. The MS_DATA bit in <i>mflag</i> must be set. Mounting a pcfs file system requires a pointer to a structure containing mount flags and local timezone information, <i>*pc_argp</i>:</p> <pre>struct pcfs_args { int timezone; /* seconds west of Greenwich */ int daylight; /* type of dst correction */ int flags; };</pre> <p>The information required in the <i>timezone</i> and <i>daylight</i> members of this structure is described in <i>ctime(3C)</i>. <i>flags</i> can contain the PCFS_MNT_FOLD_CASE flag. Fold names read from the file system to lowercase.</p>
Mounting File Systems	<p>Use the following command to mount pcfs from diskette:</p> <pre>mount -F pcfs device-special directory-name</pre> <p>You can use:</p> <pre>mount directory-name</pre> <p>if the following line is in your <i>/etc/vfstab</i> file:</p> <pre>device-special - directory-name pcfs - no rw</pre> <p>Use the following command to mount pcfs from non-diskette media:</p> <pre>mount -F pcfs device-special:logical-drive directory-name</pre> <p>You can use:</p> <pre>mount directory-name</pre> <p>if the following line is in your <i>/etc/vfstab</i> file:</p> <pre>device-special:logical_drive - directory-name pcfs - no rw</pre>

device-special specifies the special block device file for the diskette (`/dev/disketteN`) or the entire hard disk (`/dev/dsk/cNtNdNp0` for a SCSI disk, and `/dev/dsk/cNdNp0` for IDE disks) or the PCMCIA pseudo-floppy memory card (`/dev/dsk/cNtNdNsN`).

logical-drive specifies either the DOS logical drive letter (c through z) or a drive number (1 through 24). Drive letter c is equivalent to drive number 1 and represents the Primary DOS partition on the disk; drive letters d through z are equivalent to drive numbers 2 through 24, and represent DOS drives within the Extended DOS partition. Note that *device-special* and *logical-drive* must be separated by a colon.

directory-name specifies the location where the file system is mounted.

For example, to mount the Primary DOS partition from a SCSI hard disk, use:

```
mount -F pcfs /dev/dsk/cNtNdNp0:c /pcfs/c
```

To mount the first logical drive in the Extended DOS partition from an IDE hard disk, use:

```
mount -F pcfs /dev/dsk/cNdNp0:d /pcfs/d
```

To mount a DOS diskette in the first floppy drive when Volume Management is not running (see `vol(1M)`) use:

```
mount -F pcfs /dev/diskette /pcfs/a
```

If Volume Management is running, run `volcheck(1)` to automatically mount the floppy and some removable disks.

To mount a PCMCIA pseudo-floppy memory card, with Volume Management not running (or not managing the PCMCIA media), use:

```
mount -F pcfs /dev/dsk/cNtNdNsN /pcfs
```

Conventions

Files and directories created through `pcfs` must comply with either the DOS short file name convention or the long file name convention introduced with Windows 95. The DOS short file name convention is of the form *filename[.ext]*, where *filename* generally consists of from one to eight upper-case characters, while the optional *ext* consists of from one to three upper-case characters.

The long file name convention is much closer to Solaris file names. A long file name can consist of any characters valid in a short file name, lowercase letters, non-leading spaces, the characters `+, ; = []`, any number of periods, and can be up to 255 characters long. Long file names have an associated short file name for systems that do not support long file names (including earlier releases of Solaris). The short file name is not visible if the system recognizes long file names. `pcfs` generates a unique short name automatically when creating a long file name.

pcfs(7FS)

Given a long file name such as `This is a really long filename.TXT`, the short file name will generally be of the form `THISIS~N.TXT`, where *N* is a number. The long file name will probably get the short name `THISIS~1.TXT`, or `THISIS~2.TXT` if `THISIS~1.TXT` already exists (or `THISIS~3.TXT` if both exist, and so forth). If you use `pcfs` file systems on systems that do not support long file names, you may want to continue following the short file name conventions. See `EXAMPLES`.

When creating a file name, `pcfs` creates a short file name if it fits the DOS short file name format, otherwise it creates a long file name. This is because long file names take more directory space. Because the root directory of a `pcfs` file system is fixed size, long file names in the root directory should be avoided if possible.

When displaying file names, `pcfs` shows them exactly as they are on the media. This means that short names are displayed as uppercase and long file names retain their case. Earlier versions of `pcfs` folded all names to lowercase, which can be forced with the `PCFS_MNT_FOLDCASE` mount option. All file name searches within `pcfs`, however, are treated as if they were uppercase, so `readme.txt` and `ReAdMe.TXT` refer to the same file.

To format a diskette or a PCMCIA pseudo-floppy memory card in DOS format in the SunOS system, use either the `fdformat -d` or the DOS `FORMAT` command.

Boot Partitions

On x86 systems, hard drives may contain an `fdisk` partition reserved for the Solaris boot utilities. These partitions are special instances of `pcfs`. You can mount an x86 boot partition with the command:

```
mount -F pcfs device-special:boot directory-name
```

or you can use:

```
mount directory-name
```

if the following line is in your `/etc/vfstab` file:

```
device-special:boot - directory-name pcfs - no rw
```

device-special specifies the special block device file for the entire hard disk (`/dev/dsk/cNtNdNp0`)

directory-name specifies the location where the file system is mounted.

All files on a boot partition are owned by super-user. Only the super-user may create, delete, or modify files on a boot partition.

EXAMPLES

EXAMPLE 1 Sample Displays of File Names

If you copy a file `financial.data` from a UNIX file system to `pcfs`, it displays as `financial.data` in `pcfs`, but may show up as `FINANC~1.DAT` in systems that do not support long file names.

The following are legal long file names. They are also *illegal* short file names:

EXAMPLE 1 Sample Displays of File Names (Continued)

```
test.sh.orig
data+
.login
```

Other systems that do not support long file names may see:

```
TESTSH~1.ORI
DATA~1
LOGIN~1
```

The short file name is generated from the initial characters of the long file name, so differentiate names in the first few characters. For example, these names:

```
WorkReport.January.Data
WorkReport.February.Data
WorkReport.March.Data
```

result in these short names, which are not distinguishable:

```
WORKRE~1.DAT
WORKRE~2.DAT
WORKRE~13.DAT
```

These names, however:

```
January.WorkReport.Data
February.WorkReport.Data
March.WorkReport.Data
```

result in the more descriptive short names:

```
JANUAR~1.DAT
FEBRUA~1.DAT
MARCHW~1.DAT
```

FILES /usr/lib/fs/pcfs/mount pcfs mount command
/usr/kernel/fs/pcfs 32-bit kernel module

SEE ALSO chgrp(1), chown(1), dos2unix(1), eject(1), fdformat(1), unix2dos(1), volcheck(1), mount(1M), mount_pcfs(1M), vold(1M), ctime(3C), vfstab(4), pcmem(7D)

pcfs(7FS)

- WARNINGS** Do not physically eject a DOS floppy while the device is mounted as `pcfs`. If Volume Management is managing a device, use the `eject(1)` command before physically removing media.
- When mounting `pcfs` on a hard disk, make sure the first block on that device contains a valid `fdisk` partition table.
- Because `pcfs` has no provision for handling owner-IDs or group-IDs on files, `chown(1)` or `chgrp(1)` may generate various errors. This is a limitation of `pcfs`, but it should not cause problems other than error messages.
- NOTES** Only the following characters are allowed in `pcfs` short file names and extensions:
- 0-9
A-Z
\$#&!%() - { } < > ` ^ ~ | ' ,
- SunOS and DOS use different character sets and have different requirements for the text file format. Use the `dos2unix(1)` and `unix2dos(1)` commands to convert files between them.
- `pcfs` offers a convenient transportation vehicle for files between Sun workstations and PCs. Because the DOS disk format was designed for use under DOS, it does not operate efficiently under the SunOS system and should not be used as the format for a regular local storage. Instead, use `ufs` for local storage within the SunOS system.
- Although long file names can contain spaces (just as in UNIX file names), some utilities may be confused by them.
- This implementation of `pcfs` conforms to the behavior exhibited by Windows 95 version 4.00.950.
- BUGS** `pcfs` should handle the disk change condition in the same way that DOS does, so you do not need to unmount the file system to change floppies.
- When listing or searching a directory, `pcfs` does not include files with the *hidden* or *system* bits set.

NAME	pcic – Intel i82365SL PC Card Interface Controller						
DESCRIPTION	<p>The Intel i82365SL PC Card interface controller provides one or more PCMCIA PC card sockets. The <code>pcic</code> driver implements a PCMCIA bus nexus driver.</p> <p>The driver provides basic support for the Intel 82365SL and compatible chips. The chips that have been tested are:</p> <ul style="list-style-type: none"> ■ Intel 82365SL ■ Cirrus Logic PD6710/PD6720/PD6722 ■ Vadem VG365/VG465/VG468/VG469 ■ Toshiba PCIC and ToPIC ■ Ricoh RF5C366 ■ Texas Instruments PCI1130/PCI1131/PCI1031 <p>While most systems using one of the above chips will work, some systems will not be supported due to hardware designs options that may not be software detectable. Note that systems with <code>CardBus</code> interfaces are only supported in the non-legacy mode. Systems that only initialize the bridge to legacy mode and do not configure the PCI memory will not be supported.</p> <p>Direct access to the PCMCIA hardware is not supported. All device access must be through the Card Services interface of the DDI.</p>						
PRECONFIGURE	<ul style="list-style-type: none"> ■ Install your add-in PC Card adapter prior to the Solaris installation. ■ Some systems have their built-in PC Card adapter disabled by default. Enable it prior to Solaris installation. ■ Requirements for a system depend on the combination of devices to be used. A typical two-socket system needs at least 8 Kbytes of address space, 16 bytes of I/O space, and three free IRQs. Following are general guidelines: <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 20%;">Address space</td> <td>At least 8 Kbytes are required with 4 Kbytes per socket in the 640K-1MB range (not necessarily contiguous); if there are three sockets, at least 12 Kbytes are needed.</td> </tr> <tr> <td>I/O space</td> <td>At least 8 and preferably 16 bytes per socket.</td> </tr> <tr> <td>IRQs</td> <td>One per socket, plus an IRQ for the <code>pcic</code> device driver itself.</td> </tr> </table>	Address space	At least 8 Kbytes are required with 4 Kbytes per socket in the 640K-1MB range (not necessarily contiguous); if there are three sockets, at least 12 Kbytes are needed.	I/O space	At least 8 and preferably 16 bytes per socket.	IRQs	One per socket, plus an IRQ for the <code>pcic</code> device driver itself.
Address space	At least 8 Kbytes are required with 4 Kbytes per socket in the 640K-1MB range (not necessarily contiguous); if there are three sockets, at least 12 Kbytes are needed.						
I/O space	At least 8 and preferably 16 bytes per socket.						
IRQs	One per socket, plus an IRQ for the <code>pcic</code> device driver itself.						
CONFIGURATION	<p>There is one driver configuration property defined in the <code>pcic.conf</code> file.</p> <pre>interrupt-priorities=11;</pre> <p style="text-align: right;">This property must be defined and must not be modified from the default value.</p>						
Initial Installation and Configuration	<p>To perform initial installation and configuration, do the following steps:</p> <ol style="list-style-type: none"> 1. Consult the Configuration Assistant for address space, I/O space, and IRQs already used by system devices. 2. Insert the PC Card adapter. 						

pcic(7D)

3. Install the Solaris software.
4. Reboot the system.

Adding PC Card Support to a Previously Installed System

To add PC card support to a previously installed system, do the following:

1. Become root.
2. Do a reconfiguration reboot to reallocate resources:

```
# touch /reconfigure
# reboot
```
3. Insert the PC Card adapter and turn on the machine.
4. Boot the system so that the PC Card device driver begins running with the new resources allocated.

Allocating IRQs

To allocate the IRQs, do the following:

1. Boot with the Configuration Assistant so you can review the resource usage.
2. Select View/Edit Devices from the Device Tasks menu, and review the list of devices to see how many IRQs are being used. (There are 16 IRQs, from 0-15. Several IRQs are already assigned. For example, IRQ 3 is reserved for the second serial port, COM2, and IRQ 7 is reserved for the parallel port.)
3. If your system has an unused COM2 or parallel port, delete the device to free the IRQ resource for a PC Card. To do this, select the serial port device using IRQ 3 or the parallel port using IRQ 7, and choose Delete Device. Next, choose Continue to return to the Device Tasks menu, then save the configuration.
4. Boot the Solaris software.

FILES	/kernel/drv/pcic	pcic driver
	/kernel/drv/pcic.conf	pcic configuration file

SEE ALSO pcmcia(4), stp4020(7D)

NAME	pckt – STREAMS Packet Mode module
SYNOPSIS	<pre>int ioctl(fd, I_PUSH, "pckt");</pre>
DESCRIPTION	<p>pckt is a STREAMS module that may be used with a pseudo terminal to packetize certain messages. The pckt module should be pushed (see I_PUSH on streamio(7I)) onto the master side of a pseudo terminal.</p> <p>Packetizing is performed by prefixing a message with an M_PROTO message. The original message type is stored in the 1 byte data portion of the M_PROTO message.</p> <p>On the read-side, only the M_PROTO, M_PCPROTO, M_STOP, M_START, M_STOPI, M_STARTI, M_IOCTL, M_DATA, M_FLUSH, and M_READ messages are packetized. All other message types are passed upstream unmodified.</p> <p>Since all unread state information is held in the master's stream head read queue, flushing of this queue is disabled.</p> <p>On the write-side, all messages are sent down unmodified.</p> <p>With this module in place, all reads from the master side of the pseudo terminal should be performed with the getmsg(2) or getpmsg() function. The control part of the message contains the message type. The data part contains the actual data associated with that message type. The onus is on the application to separate the data into its component parts.</p>
SEE ALSO	<p>getmsg(2), ioctl(2), ldterm(7M), ptem(7M), streamio(7I), termio(7I)</p> <p><i>STREAMS Programming Guide</i></p>

pcmem(7D)

NAME	pcmem – PCMCIA memory card nexus driver
DESCRIPTION	<p>The <code>pcmem</code> driver identifies the type of memory card in the system and will allow future support of other memory device types.</p> <p>The PCMCIA memory card nexus driver supports PCMCIA memory card client drivers. There are no user-configurable options for this driver.</p>
FILES	<code>/kernel/drv/pcmem</code> <code>pcmem</code> driver
SEE ALSO	pcram(7D)

NAME	pcn – AMD PCnet Ethernet controller device driver
SYNOPSIS	/dev/pcn
DESCRIPTION	<p>The pcn Ethernet driver is a multi-threaded, loadable, clonable driver for the AMD PCnet family of Ethernet controllers that use the Generic LAN Driver (GLD) facility to implement the required STREAMS and Data Link Provider (see dlpi(7P)) interfaces.</p> <p>This driver supports a number of integrated motherboards and add-in adapters based on the AMD PCnet-ISA, PCnet-PCI, and PCnet-32 controller chips. The pcn driver functions include controller initialization, frame transmit and receive, functional addresses, promiscuous and multicast support, and error recovery and reporting.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning character-special device, /dev/pcn, is used to access all PCnet devices installed in the system.</p>
pcn and DLPI	<p>The pcn driver uses the Solaris GLD module which handles all the STREAMS and DLPI specific functions of the driver. It is a <i>style 2</i> DLPI driver and therefore supports only the connectionless mode of data transfer. Thus, a DLPI user should issue a DL_ATTACH_REQ primitive to select the device to be used. Valid DLPI primitives are defined in <sys/dlpi.h>. Refer to dlpi(7P) for more information.</p>
	<p>The device is initialized on the first attach and de-initialized (stopped) on the last detach.</p>
	<p>The values returned by the driver in the DL_INFO_ACK primitive in response to a DL_INFO_REQ from the user are:</p>
	<ul style="list-style-type: none"> ■ Maximum SDU is 1500 (ETHERMTU - defined in <sys/ethernet.h>). ■ Minimum SDU is 0. ■ DLSAP address length is 8. ■ MAC type is DL_ETHER. ■ sap length value is -2, meaning the physical address component is followed immediately by a 2-byte sap component within the DLSAP address. ■ Service mode is DL_CLDLS. ■ No optional quality of service (QOS) support is included at present, accordingly, the QOS fields are 0. ■ Provider style is DL_STYLE2. ■ Version is DL_VERSION_2. ■ Broadcast address value is the Ethernet/IEEE broadcast address (FF:FF:FF:FF:FF:FF).
	<p>Once in the DL_ATTACHED state, the user must send a DL_BIND_REQ to associate a particular Service Access Point (SAP) with the stream.</p>

pcn(7D)

Known Problems and Limitations

- Occasional data corruption has occurred when `pcn` and `pcscsi` drivers in HP Vectra XU 5/90 and Compaq Deskpro XL systems are used under high network and SCSI loads. These drivers do not perform well in a production server. A possible workaround is to disable the `pcn` device with the system BIOS and use a separate add-in network interface.
- The Solaris `pcn` driver does not support IRQ 4.

FILES

`/dev/pcn` Character special device
`/kernel/drv/pcn.conf` Configuration file

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO

`attributes(5)`, `standards(5)`, `dlpi(7P)`, `streamio(7I)`

Writing Device Drivers

STREAMS Programming Guide

NAME	pcram – PCMCIA RAM memory card device driver
SYNOPSIS	memory@<socket>/pcram@<technology>, 0:c memory@<socket>/pcram@<technology>, 0:c, raw
DESCRIPTION	The PCMCIA RAM memory card device driver supports disk-like I/O access to any standard PCMCIA static random access memory (SRAM) card and dynamic random access memory (DRAM) card. The driver supports standard PCMCIA SRAM/DRAM cards that contain a Card Information Structure (CIS). RAM card densities in the 512Kilobytes to 64Mbyte range are supported.
PRECONFIGURATION	When a PC card memory device is recognized, the <code>pcram</code> device driver is automatically loaded, the physical address allocated, and special files created (if they do not already exist).
Known Problems and Limitations	The Solaris <code>pcmem</code> driver is not capable of handling "combo" memory cards with multiple types of memory on them (for example, combined SRAM and nonvolatile FLASH). Inserting such a card into a system running the Solaris software may cause a system panic. Because the PC card memory device is designed as a pseudo-floppy diskette type, the only utility that can be used for formatting is <code>fdformat(1)</code> .
CONFIGURATION	Configuration topics include initial installation and configuration, identifying an unrecognized device, special files, using PC card memory devices, and hot-plugging.
Initial Installation and Configuration	<ol style="list-style-type: none"> 1. Install the Solaris software. 2. Boot the system. 3. Insert the card.
Identifying an Unrecognized Device	<p>If you insert a memory device and it is not recognized (no special files created), use the <code>prtconf</code> command.</p> <ol style="list-style-type: none"> 1. Become root. 2. Run the <code>prtconf -D</code> command to display the configuration recognized by the system. A recognized device will appear in the <code>prtconf</code> output. For example: <pre># prtconf -D . . . pcic, instance #0 (driver name: pcic) . . . memory, instance #0 (driver name: pmem) pcram, instance #0 (driver name: pcram)</pre> 3. If your memory device does not appear at the end of the <code>prtconf</code> output, it is not supported and cannot be used with the <code>pcram</code> driver.
Special Files	<p>The special files created for PC card memory devices act like disks and have names in the form <code>/dev/dsk/c#t#d#p#</code> or <code>/dev/dsk/c#t#d#s#</code>. Abbreviations used in the names are:</p> <ul style="list-style-type: none"> ■ <code>c#</code> Controller #

pcram(7D)

- `t#` Card technology type #, defined as follows:

0 Null--no device
1 ROM
2 OTPROM (One Time PROM)
3 UV EPROM
4 EEPROM
5 Flash EPROM
6 SRAM
7 DRAM

- `d#` Device region of type #, usually zero
- `p#` fdisk partition #
- `s#` Solaris slice #

Note – A device name can be specified either by a partition name (`p#`) or a slice name (`s#`), but not both.

Using PC Card Memory Devices

Since the Solaris Volume Management software recognizes PC Card memory devices, no special `vold` configuration is required. If you do not want to use `vold` to manage your PC card memory devices, comment out the `use pcmem` line in the `/etc/vold.conf` file. To comment out a line, insert a `#` character at the beginning of the line.

PC Card memory devices do not need to have file systems on them, though typically, before using a new PC Card memory card, you will want to create a file system on it. DOS PCFS is the best format to use. (You can use virtually any file system format on a PC card memory card, but most other file system formats are platform-dependent, making them unsuitable for moving data between different types of machines. See "Using a PCMCIA Memory Card" in the *OpenWindows Advanced User's Guide*.)

Note – If you want to redirect the output of a `tar` command (or `dd` or `cpio`) to a PC card memory device, first create a file system on the card, using the `fdformat(1)` command without arguments. The card must be reformatted before it can be written on again.

Hot-Plugging

If a memory card is removed while in use, the device driver returns errors until the memory card is inserted into the appropriate socket. Close and reopen the device with the card reinserted, and the memory card will work.

- If you remove the card while in use as a file system, unmount the file system using the `umount` command. Then reinsert the card and remount the file system using the `mount` command.
- If you remove the card and interrupt a `tar` or `cpio` process, stop the process, reinsert the card, and restart the process.

FILES	<code>/kernel/drv/pcram</code>	pcram driver
	<code>/dev/dsk/cntndnsn</code>	block files

`/dev/rdisk/cntndnsn`

raw files where:

`cn` controller *n*

`tn` technology type *n*

`0x1ROM,0x2OTEPROM,0x3EPROM,`

`0x4EEPROM,0x5FLASH,0x6SRAM,`

`0x7DRAM`

`dn` technology region in type *n*

`sn` slice *n*

SEE ALSO `fdformat(1)`, `pcmcia(4)`, `dkio(7I)`, `pcmem(7D)`

pcscsi(7D)

NAME pcscsi – low-level module for the AMD PCscsi, PCscsi II, PCnet-SCSI, and Qlogic QLA510 PCI-to-SCSI bus adapters

SYNOPSIS `pcscsi@ioaddr, 0`

DESCRIPTION The `pcscsi` module provides low-level interface functions between the common disk/tape I/O subsystem and the Am53C974 (PCscsi), Am53C974A (PCscsi II), Am79C974 (PCnet-SCSI) (SCSI device only), and the Qlogic QLA510 SCSI controllers.

The `pcscsi` module can be configured for disk and streaming tape support for one host bus adapter device. Each host bus adapter device must be the sole initiator on a SCSI bus. Auto-configuration code determines if the adapter is present on the PCI bus, what its configuration is, and what types of devices are attached to it.

For PCI devices, configuration is done through the PCI BIOS. Configuration settings can be accessed through a CMOS utility.

PRECONFIGURE The net component of the PCnet-SCSI host bus adapter requires the Solaris [pcn\(7D\)](#) driver. See *AMD PCnet Ethernet (PCnet-PCI, PCnet-PCI II, PCnet-Fast)* for information on Ethernet configuration capabilities.

Known Problems and Limitations Occasional data corruption has occurred when `pcn` and `pcscsi` drivers in HP Vectra XU 5/90 and Compaq Deskpro XL systems are used under high network and SCSI loads. These drivers do not perform well in a production server. A possible workaround is to disable the `pcn` device with the system BIOS and use a separate add-in network interface.

The SCSI tagged queuing option is not supported.

CONFIGURATION The driver attempts to initialize itself in accordance with the PCI BIOS configuration settings..

There are no user-configurable options; information found in the `pcscsi.conf` configuration file is used by the I/O subsystem only.

FILES `/kernel/drv/pcscsi.conf` Configuration file

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO [driver.conf\(4\)](#), [sysbus\(4\)](#), [attributes\(5\)](#)

NAME	pcser – PCMCIA serial card device driver
SYNOPSIS	<pre>serial@<socket>:pcser serial@<socket>:pcser, cu</pre>
DESCRIPTION	The <code>pcser</code> PCMCIA serial card device driver supports asynchronous serial I/O access to any PCMCIA card that complies with Revision 2.1 of the PCMCIA standard and which represents an 8250-type UART interface.
PRECONFIGURE	If a PC card modem or serial device is recognized, the <code>pcser</code> device driver is automatically loaded, ports and IRQs allocated, and special files created (if they don't already exist).
CONFIGURATION	Configuration steps include initial installation and configuration, identifying an unrecognized device and misidentifying a recognized device.
Initial Installation and Configuration	<ol style="list-style-type: none"> 1. Install the Solaris software. 2. Boot the system. 3. Insert the modem or serial device.
Identifying an Unrecognized Device	<p>If you insert a PC card modem or serial device and it is not recognized (that is, no special files are created under <code>/dev/cua</code> or <code>/dev/term</code>), use the <code>prtconf</code> command to find the problem:</p> <ol style="list-style-type: none"> 1. Become root. 2. Run the <code>prtconf -D</code> command to see if your modem or serial device is recognized. An unrecognized device will appear at the end of the <code>prtconf</code> output. For example: <pre># prtconf -D . . . pcic, instance #0 (driver name: pcic) . . . pccard111.222 (driver not attached)</pre> 3. If your device is not recognized, use the <code>add_drv</code> command to add the name of your device as another known alias for <code>pcser</code> devices. For example, type the following at the command line: <pre># add_drv -i "pccard111.222" pcser</pre> <p>Note – Include the double quotes in single quotes to keep the shell from stripping out the double quotes. Use the identification string listed in the <code>prtconf</code> output. Use the entire string in the <code>add_drv</code> command. See <code>add_drv(1M)</code>.</p>
Misidentifying a Recognized Device	<ol style="list-style-type: none"> 1. Run the <code>prtconf -D</code> command to see if your modem or serial device is erroneously recognized as a memory card. If the device is incorrectly recognized as a memory card, the output of the <code>prtconf</code> command could show: <pre># prtconf -D . . . pcic, instance #0 (driver name: pcic) . . .</pre>

pcser(7D)

```
memory, instance #0 (driver name: pcmem)
pcram, instance #0 (driver name: pcram)
```

2. Use the Configuration Assistant to identify the memory resource conflict, and add correct information for the device on the View/Edit Devices menu. Typically, the problem may be a resource conflict between device memory settings. A PC Card adapter chip that is not fully supported may also be the cause of the problem.
3. To work properly with the Solaris operating environment, all devices must be accounted for, even those the Solaris environment does not support. The Configuration Assistant software accounts for all devices in your system.

Additional Configuration

When adding a new serial port or modem to the system, you often need to edit configuration files so that applications can use the new communications port. For example, the `/etc/uucp/Devices` file needs to be updated to use UUCP. See *Overview of UUCP* in the *System Administration Guide*. For PPP on the serial port, see `pppd(1M)` and *Solaris PPP Overview* in the *System Administration Guide*.

Special Files

The serial devices in `/dev/term` and `/dev/cua` are named by socket number. A card inserted in socket 0 is `pc0`, and socket 1 is `pc1`.

Hot Plugging

If a PC Card modem or serial device is unplugged while in use, the device driver returns errors until the card is replaced in the socket.

The device must be closed and reopened with the card reinserted before the device begins working again. The restart process depends on the application. For example, a `tip` session automatically exits when a card in use is unplugged. To restart the system, you must restart the `tip` session.

FILES

<code>/kernel/drv/pcser</code>	pcser driver
<code>/dev/term/pcn</code>	dial-in devices
<code>/dev/cua/pcn</code>	dial-out devices where: <i>n</i> is the PCMCIA physical socket number.

SEE ALSO

`cu(1C)`, `tip(1)`, `uucp(1C)`, `autopush(1M)`, `pcmcia(1M)`, `ports(1M)`, `ioctl(2)`, `open(2)`, `pcmcia(4)`, `termio(7I)`, `ldterm(7M)`, `ttcompat(7M)`

DIAGNOSTICS

`pcser: socket n soft silo overflow`
The driver's character input ring buffer overflowed before it could be serviced.

`pcser: socket n unable to get CIS information`
The CIS on the card has incorrect information or is in an incorrect format. This message usually indicates a non-compliant card.

NAME	pfb – Sun XVR-100 Graphics Accelerator device driver
DESCRIPTION	The pfb driver is the device driver for the Sun XVR-100 Graphics Accelerator.
FILES	<code>/dev/fbs/pfb\fin</code> device special file for XVR-100 single screen <code>/dev/fbs/pfb\fin\fPa</code> device special file for the XVR-100 first video out <code>/dev/fbs/pfb\fin\fPb</code> device special file for the XVR-100 second video out
SEE ALSO	SUNWpfb_config(1M)

pf_key(7P)

NAME	pf_key – security association database
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/socket.h> #include <net/pfkeyv2.h> int socket (PF_KEY, SOCK_RAW, PF_KEY_V2);</pre>
DESCRIPTION	<p>Keying information for IPsec security services is maintained in security association databases (SADBs). The security associations (SAs) are used to protect both inbound and outbound packets.</p> <p>A user process (or possibly multiple co-operating processes) maintains SADBs by sending messages over a special kind of socket. This is analogous to the method described in route(7P). Only a superuser may access an SADB.</p> <p>SunOS applications that use PF_KEY include ipseckey(1M) and in.iked(1M).</p> <p>The operating system may spontaneously emit messages in response to external events, such as a request for a new SA for an outbound datagram, or to report the expiration of an existing SA.</p> <p>One opens the channel for passing SADB control messages by using the socket call shown in the SYNOPSIS section above. More than one key socket can be open per system.</p> <p>Messages are formed by a small base header, followed by a number, zero or more, of extension messages, some of which require additional data following them. The base message and all extensions must be eight-byte aligned. An example message is the GET message, which requires the base header, the SA extension, and the ADDRESS_DST extension.</p>
Messages	<p>Messages include:</p> <pre>#define SADB_GETSPI /* Get a new SPI value from the system. */ #define SADB_UPDATE /* Update an SA. */ #define SADB_ADD /* Add a fully-formed SA. */ #define SADB_DELETE /* Delete an SA. */ #define SADB_GET /* Get an SA */ #define SADB_ACQUIRE /* Kernel needs a new SA. */ #define SADB_REGISTER /* Register to receive ACQUIRE messages. */ #define SADB_EXPIRE /* SA has expired. */ #define SADB_FLUSH /* Flush all SAs. */ #define SADB_DUMP /* Get all SAs. (Unreliable) */ #define SADB_X_PROMISC /* Listen promiscuously */ #define SADB_X_INVERSE_ACQUIRE /* Query kernel policy, get an ACQUIRE in return. */</pre> <p>The base message header consists of:</p> <pre>struct sadb_msg { uint8_t sadb_msg_version; /* Set to PF_KEY_V2, for compatibility */ uint8_t sadb_msg_type; /* Message type */ uint8_t sadb_msg_errno; /* Why message failed */</pre>


```

uint8_t  sadb_msg_satype; /* Which security service */
uint16_t sadb_msg_len;   /* Length in 8-byte units */
uint16_t sadb_msg_reserved; /* Zero out */
#define sadb_x_msg_diagnostic sadb_msg_reserved
                               /* Extended diagnostics for errors */
uint32_t sadb_msg_seq;   /* For message originator */
uint32_t sadb_msg_pid;   /* Identify originator */
};

```

Extension types include:

```

#define SADB_EXT_SA /* SA information */
#define SADB_EXT_LIFETIME_HARD /* Hard lifetime */
#define SADB_EXT_LIFETIME_SOFT /* Soft lifetime */
#define SADB_EXT_ADDRESS_SRC /* Source address */
#define SADB_EXT_ADDRESS_DST /* Destination address */
#define SADB_EXT_ADDRESS_PROXY /* Proxy address */
#define SADB_EXT_KEY_AUTH /* Authentication key */
#define SADB_EXT_KEY_ENCRYPT /* Encryption key */
#define SADB_EXT_IDENTITY_SRC /* Source certificate ID */
#define SADB_EXT_IDENTITY_DST /* Destination certificate ID */
#define SADB_EXT_SENSITIVITY /* Sensitivity information */
#define SADB_EXT_PROPOSAL /* Security proposal */
#define SADB_EXT_SUPPORTED_AUTH /* Supported authentication algorithms */
#define SADB_EXT_SUPPORTED_ENCRYPT /* Supported encryption algorithms */
#define SADB_EXT_SPIRANGE /* Range of possible SPIs */
#define SADB_X_EXT_EREG /* Register for extended ACQUIRE */
#define SADB_X_EXT_EPROP /* Extended ACQUIRE proposals */
#define SADB_X_EXT_KM_COOKIE /* Indicates which KM derived SA. */

```

Extension headers include:

Generic Extension Header

```

struct sadb_ext {
    uint16_t sadb_ext_len; /* In 64-bit words, inclusive */
    uint16_t sadb_ext_type; /* 0 is reserved */
};

```

Security Association Information Extension

```

struct sadb_sa {
    uint16_t sadb_sa_len;
    uint16_t sadb_sa_exttype; /* ASSOCIATION */
    uint32_t sadb_sa_spi;
    uint8_t sadb_sa_replay;
    uint8_t sadb_sa_state;
    uint8_t sadb_sa_auth;
    uint8_t sadb_sa_encrypt;
    uint32_t sadb_sa_flags;
};

```

Lifetime Extension

pf_key(7P)

```
struct sadb_lifetime {
    uint16_t sadb_lifetime_len;
    uint16_t sadb_lifetime_exttype;    /* SOFT, HARD, CURRENT */
    uint32_t sadb_lifetime_allocations;
    uint64_t sadb_lifetime_bytes;
    uint64_t sadb_lifetime_addtime;
    uint64_t sadb_lifetime_usetime;
};
```

Address Extension

```
struct sadb_address {
    uint16_t sadb_address_len;
    uint16_t sadb_address_exttype;    /* SRC, DST, PROXY */
    uint8_t sadb_address_proto;       /* Proto for ports... */
    uint8_t sadb_address_prefixlen;   /* Prefix length. */
    uint16_t sadb_address_reserved;   /* Padding */
                                        /* Followed by a sockaddr structure.*/
};
```

Keying Material Extension

```
struct sadb_key {
    uint16_t sadb_key_len;
    uint16_t sadb_key_exttype;        /* AUTH, ENCRYPT */
    uint16_t sadb_key_bits;
    uint16_t sadb_key_reserved;
    /* Followed by actual key(s) in canonical (outbound proc.) order. */
};
```

Identity Extension

```
struct sadb_ident {
    uint16_t sadb_ident_len;
    uint16_t sadb_ident_exttype;     /* SRC, DST, PROXY */
    uint16_t sadb_ident_type;        /* FQDN, USER_FQDN, etc. */
    uint16_t sadb_ident_reserved;    /* Padding */
    uint64_t sadb_ident_id;          /* For userid, etc. */
    /* Followed by an identity null-terminate C string if present. */
};
```

Sensitivity/Integrity Extension

```
struct sadb_sens {
    uint16_t sadb_sens_len;
    uint16_t sadb_sens_exttype;     /* SENSITIVITY */
    uint32_t sadb_sens_dpd;
    uint8_t sadb_sens_sens_level;
    uint8_t sadb_sens_sens_len;     /* 64-bit words */
    uint8_t sadb_sens_integ_level;
    uint8_t sadb_sens_integ_len;    /* 64-bit words */
    uint32_t sadb_sens_reserved;
    /*
     * followed by two uint64_t arrays
    */
};
```

```

        * uint64_t sadb_sens_bitmap[sens_bitmap_len];
        * uint64_t integ_bitmap[integ_bitmap_len];
        */
};

```

Proposal Extension

```

struct sadb_prop {
    uint16_t sadb_prop_len;
    uint16_t sadb_prop_len;
    uint16_t sadb_prop_exttype;      /* PROPOSAL, X_EPROP */
    union {
        struct {
            uint8_t sadb_prop_lenres_replay;
            uint8_t sadb_prop_lenres_eres;
            uint16_t sadb_prop_lenres_numcombs;
        } sadb_prop_lenres;
        struct {
            uint8_t sadb_prop_oldres_replay;
            uint8_t sadb_prop_oldres_reserved[3];
        } sadb_prop_oldres;
    } sadb_prop_u;
#define sadb_prop_replay sadb_prop_u.sadb_prop_lenres.sadb_prop_lenres_replay
#define sadb_prop_reserved \

    sadb_prop_u.sadb_prop_oldres.sadb_prop_oldres_reserved
#define sadb_x_prop_ereserved \

    sadb_prop_u.sadb_prop_lenres.sadb_prop_lenres_eres
#define sadb_x_prop_numcombs \

    sadb_prop_u.sadb_prop_lenres.sadb_prop_lenres_numcombs \

/* Followed by sadb_comb[] array or sadb_ecomb[] array. */
};
};

```

Combination Instance for a Proposal

```

struct sadb_comb {
    uint8_t sadb_comb_auth;
    uint8_t sadb_comb_encrypt;
    uint16_t sadb_comb_flags;
    uint16_t sadb_comb_auth_minbits;
    uint16_t sadb_comb_auth_maxbits;
    uint16_t sadb_comb_encrypt_minbits;
    uint16_t sadb_comb_encrypt_maxbits;
    uint32_t sadb_comb_reserved;
    uint32_t sadb_comb_soft_allocations;
    uint32_t sadb_comb_hard_allocations;
    uint64_t sadb_comb_soft_bytes;
    uint64_t sadb_comb_hard_bytes;
    uint64_t sadb_comb_soft_addtime;
    uint64_t sadb_comb_hard_addtime;
    uint64_t sadb_comb_soft_usetime;
    uint64_t sadb_comb_hard_usetime;
};

```

```
};
```

Extended Combination

```
struct sadb_x_ecomb {
    uint8_t sadb_x_ecomb_numalgs;
    uint8_t sadb_x_ecomb_reserved;
    uint16_t sadb_x_ecomb_flags; /* E.g. PFS? */
    uint32_t sadb_x_ecomb_reserved2;
    uint32_t sadb_x_ecomb_soft_allocations;
    uint32_t sadb_x_ecomb_hard_allocations;
    uint64_t sadb_x_ecomb_soft_bytes;
    uint64_t sadb_x_ecomb_hard_bytes;
    uint64_t sadb_x_ecomb_soft_addtime;
    uint64_t sadb_x_ecomb_hard_addtime;
    uint64_t sadb_x_ecomb_soft_usetime;
    uint64_t sadb_x_ecomb_hard_usetime;
};
```

Extended Combination Algorithm Descriptors

```
struct sadb_x_algdesc {
    uint8_t sadb_x_algdesc_satype; /* ESP, AH, etc. */
    uint8_t sadb_x_algdesc_algtype; /* AUTH, CRYPT, COMPRESS */
    uint8_t sadb_x_algdesc_alg; /* DES, 3DES, MD5, etc. */
    uint8_t sadb_x_algdesc_reserved;
    uint16_t sadb_x_algdesc_minbits; /* Bit strengths. */
    uint16_t sadb_x_algdesc_maxbits;
};
```

Extended Register

```
struct sadb_x_ereg {
    uint16_t sadb_x_ereg_len;
    uint16_t sadb_x_ereg_exttype; /* X_EREG */
    uint8_t sadb_x_ereg_satypes[4]; /* Array of SA types, 0-terminated.
};
```

Key Management Cookie

```
struct sadb_x_kmc {
    uint16_t sadb_x_kmc_len;
    uint16_t sadb_x_kmc_exttype; /* X_KM_COOKIE */
    uint32_t sadb_x_kmc_proto; /* KM protocol */
    uint32_t sadb_x_kmc_cookie; /* KMP-specific */
    uint32_t sadb_x_kmc_reserved; /* Reserved; must be zero */
};
```

Supported Algorithms Extension

```
struct sadb_supported {
    uint16_t sadb_supported_len;
    uint16_t sadb_supported_exttype;
};
```

```
uint32_t sadb_supported_reserved;
};
```

Algorithm Instance

```
struct sadb_alg {
    uint8_t sadb_alg_id;          /* Algorithm type. */
    uint8_t sadb_alg_ivlen;      /* IV len, in bits */
    uint16_t sadb_alg_minbits;   /* Min. key len (in bits) */
    uint16_t sadb_alg_maxbits;   /* Max. key length */
    uint16_t sadb_alg_reserved;
};
```

SPI Extension Range

```
struct sadb_spirange {
    uint16_t sadb_spirange_len;
    uint16_t sadb_spirange_exttype; /* SPI_RANGE */
    uint32_t sadb_spirange_min
    uint32_t sadb_spirange_max;
    uint32_t sadb_spirange_reserved;
};
```

Message Use and Behavior

Each message has a behavior. A behavior is defined as where the initial message travels, for example, user to kernel, and what subsequent actions are expected to take place. Contents of messages are illustrated as:

```
<base, REQUIRED EXTENSION, REQ., (OPTIONAL EXTENSION), (OPT)>
```

The SA extension is sometimes used only for its SPI field. If all other fields must be ignored, this is represented by SA(*).

The lifetime extensions are represented with one to three letters after the word lifetime, representing (H)ARD, (S)OFT, and (C)URRENT.

The address extensions are represented with one to three letters after the word "address," representing (S)RC, (D)ST, (P)ROXY.

Note that when an error occurs, only the base header is sent. Typical errors include:

EINVAL	Various message improprieties, including SPI ranges that are malformed, weak keys, and others. If EINVAL is returned, an application should look at the <code>sadb_x_msg_diagnostic</code> field of the <code>sadb_msg</code> structure. It contains one of many possible causes for EINVAL. See <code>net/pfkeyv2.h</code> for values, all of the form <code>SADB_X_DIAGNOSTIC_</code> .
ENOMEM	Needed memory was not available.
ENSGSIZ	Message exceeds the maximum length allowed.
EEXIST	SA (that is being added or created with <code>GETSPI</code>) already exists.

pf_key(7P)

ESRCH SA could not be found.

The following are examples of message use and behavior:

SADB_GETSPI

Send a SADB_GETSPI message from a user process to the kernel.

```
<base, address, SPI range>
```

The kernel returns the SADB_GETSPI message to all listening processes.

```
<base, SA(*), address (SD)>
```

SADB_UPDATE

Send a SADB_UPDATE message from a user process to the kernel.

```
<base, SA, (lifetime(HS),) address(SD), (address(P), key (AE),  
  (identity(SD),) (sensitivity)>c
```

The kernel returns the SADB_UPDATE message to all listening processes.

```
<base, SA(*), address (SD)>
```

SADB_ADD

Send a SADB_ADD message from a user process to the kernel.

```
<base, SA, (lifetime(HS),) address(SD), (address(P),) key (AE),  
  (identity(SD),) (sensitivity)>
```

The kernel returns the SADB_ADD message to all listening processes.

```
<base, SA, (lifetime(HS),) address (SD),  
  (identity (SD),) (sensitivity)>
```

SADB_DELETE

Send a SADB_DELETE message from a user process to the kernel.

```
<base, SA (*), address (SD)>
```

The kernel returns the SADB_DELETE message to all listening processes.

```
<base, SA (*), address (SD)>
```

SADB_GET

Send a SADB_GET message from a user process to the kernel.

```
<base, SA (*), address (SD)>
```

The kernel returns the SADB_GET message to the socket that sent the SADB_GET message.

```
<base, SA , (lifetime (HSC),) address (SD), (address (P),) key (AE),
  (identity (SD),) (sensitivity)>
```

SADB_ACQUIRE

The kernel sends a SADB_ACQUIRE message to registered sockets. Note that any GETSPI, ADD, or UPDATE calls in reaction to an ACQUIRE must fill in the `sadb_msg_seq` of those messages with the one in the ACQUIRE message. The address (SD) extensions must have the port fields filled in with the port numbers of the session requiring keys if appropriate.

```
<base, address (SD), (address(P)), (identity(SD),)
  (sensitivity,) proposal>
```

Extended ACQUIRE will have a slightly different format. The `sadb_msg_satype` field is 0, and the extension contains the desired combination(s) of security protocols.

```
<base, address (SD), (address(P)), (identity(SD),)
  (sensitivity,) eprop>
```

If key management fails, send an SADB_ACQUIRE to indicate failure.

```
<base>
```

SADB_X_INVERSE_ACQUIRE

For inbound Key Management processing, a Key Management application may wish to consult the kernel for its policy. The application should send to the kernel:

```
<base, address (SD)>
```

The kernel returns a message similar to a kernel-generated extended ACQUIRE:

```
<base, address (SD), (address(P)), (identity(SD),)
  (sensitivity,) eprop>
```

SADB_REGISTER

Send a SADB_REGISTER message from a user process to the kernel.

```
<base>
```

pf_key(7P)

The kernel returns the `SADB_REGISTER` message to registered sockets, with algorithm types supported by the kernel being indicated in the supported algorithms field. Note that this message may arrive asynchronously due to an algorithm being loaded or unloaded into a dynamically linked kernel.

```
<base, supported>
```

There is also the extended `REGISTER`, which will allow this process to receive extended `ACQUIRES`.

```
<base, ereg>
```

Which returns a series of `SADB_REGISTER` replies (one for each security protocol registered) from the kernel.

SADB_EXPIRE

The kernel sends a `SADB_EXPIRE` message to all listeners when the soft limit of a security association has been expired.

```
<base, SA, lifetime (C and one of HS), address (SD)>
```

SADB_FLUSH

Send a `SADB_FLUSH` message from a user process to the kernel.

```
<base>
```

The kernel returns the `SADB_FLUSH` message to all listening sockets.

```
<base>
```

SADB_DUMP

Send a `SADB_DUMP` message from a user process to the kernel.

```
<base>
```

Several `SADB_DUMP` messages will return from the kernel to the sending socket.

```
<base, SA, (lifetime (HSC),) address (SD), (address (P),) key (AE),  
  (identity (SD),) sensitivity)>
```

To mark the end of a dump a single base header arrives with its `sadb_mdg_seq` set to 0.

```
<base>
```


SADB_X_PROMISC

Send a SADB_X_PROMISC message from a user process to the kernel.

<base>

The kernel returns the SADB_X_PROMISC message to all listening processes.

<base>

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr
Interface Stability	Evolving

SEE ALSO `in.iked(1M)`, `ipseckey(1M)`, `ipsec(7P)`, `ipsecah(7P)`, `ipsecesp(7P)`, `route(7P)`

McDonald, D.L., Metz, C.W., and Phan, B.G., *RFC 2367, PF_KEY Key Management API, Version 2*, The Internet Society, July 1998.

NOTES Time-based lifetimes may not expire with exact precision in seconds because kernel load may affect the aging of SA's.

pfmod(7M)

NAME	pfmod – STREAMS Packet Filter Module
SYNOPSIS	<pre>#include <sys/pfmod.h> ioctl(fd, IPUSH, "pfmod");</pre>
DESCRIPTION	<p>pfmod is a STREAMS module that subjects messages arriving on its read queue to a packet filter and passes only those messages that the filter accepts on to its upstream neighbor. Such filtering can be very useful for user-level protocol implementations and for networking monitoring programs that wish to view only specific types of events.</p>
Read-side Behavior	<p>pfmod applies the current packet filter to all M_DATA and M_PROTO messages arriving on its read queue. The module prepares these messages for examination by first skipping over all leading M_PROTO message blocks to arrive at the beginning of the message's data portion. If there is no data portion, pfmod accepts the message and passes it along to its upstream neighbor. Otherwise, the module ensures that the part of the message's data that the packet filter might examine lies in contiguous memory, calling the pullupmsg(9F) utility routine if necessary to force contiguity. (Note: this action destroys any sharing relationships that the subject message might have had with other messages.) Finally, it applies the packet filter to the message's data, passing the entire message upstream to the next module if the filter accepts, and discarding the message otherwise. See PACKET FILTERS below for details on how the filter works.</p> <p>If there is no packet filter yet in effect, the module acts as if the filter exists but does nothing, implying that all incoming messages are accepted. The IOCTLS section below describes how to associate a packet filter with an instance of pfmod.</p> <p>pfmod passes all other messages through unaltered to its upper neighbor.</p>
Write-side Behavior	<p>pfmod intercepts M_IOCTL messages for the <i>ioctl</i> described below. The module passes all other messages through unaltered to its lower neighbor.</p>
IOCTLS	<p>pfmod responds to the following <i>ioctl</i>.</p> <p>PFIOCSETF This <i>ioctl</i> directs the module to replace its current packet filter, if any, with the filter specified by the struct packetfilt pointer named by its final argument. This structure is defined in <code><sys/pfmod.h></code> as:</p> <pre>struct packetfilt { uchar_t Pf_Priority; /* priority of filter */ uchar_t Pf_FilterLen; /* length of filter cmd list */ ushort_t Pf_Filter[ENMAXFILTERS]; /* filter command list */ };</pre> <p>The Pf_Priority field is included only for compatibility with other packet filter implementations and is otherwise ignored. The packet filter itself is specified in the Pf_Filter array as a sequence of two-byte commands, with the Pf_FilterLen field giving the number of commands in the sequence. This implementation restricts the maximum number of commands in a filter (ENMAXFILTERS) to 255. The next section describes the available commands and their semantics.</p>

PACKET FILTERS

A packet filter consists of the filter command list length (in units of `ushort_ts`), and the filter command list itself. (The priority field mentioned above is ignored in this implementation.) Each filter command list specifies a sequence of actions that operate on an internal stack of `ushort_ts` (“shortwords”). Each shortword of the command list specifies one of the actions `ENF_PUSHLIT`, `ENF_PUSHZERO`, `ENF_PUSHONE`, `ENF_PUSHFFFF`, `ENF_PUSHFF00`, `ENF_PUSH00FF`, or `ENF_PUSHPWORD+n`, which respectively push the next shortword of the command list, zero, one, `0xFFFF`, `0xFF00`, `0x00FF`, or shortword *n* of the subject message on the stack, and a binary operator from the set {`ENF_EQ`, `ENF_NEQ`, `ENF_LT`, `ENF_LE`, `ENF_GT`, `ENF_GE`, `ENF_AND`, `ENF_OR`, `ENF_XOR`} which then operates on the top two elements of the stack and replaces them with its result. When both an action and operator are specified in the same shortword, the action is performed followed by the operation.

The binary operator can also be from the set {`ENF_COR`, `ENF_CAND`, `ENF_CNOR`, `ENF_CNAND`}. These are “short-circuit” operators, in that they terminate the execution of the filter immediately if the condition they are checking for is found, and continue otherwise. All pop two elements from the stack and compare them for equality; `ENF_CAND` returns false if the result is false; `ENF_COR` returns true if the result is true; `ENF_CNAND` returns true if the result is false; `ENF_CNOR` returns false if the result is true. Unlike the other binary operators, these four do not leave a result on the stack, even if they continue.

The short-circuit operators should be used when possible, to reduce the amount of time spent evaluating filters. When they are used, you should also arrange the order of the tests so that the filter will succeed or fail as soon as possible; for example, checking the IP destination field of a UDP packet is more likely to indicate failure than the packet type field.

The special action `ENF_NOPUSH` and the special operator `ENF_NOP` can be used to only perform the binary operation or to only push a value on the stack. Since both are (conveniently) defined to be zero, indicating only an action actually specifies the action followed by `ENF_NOP`, and indicating only an operation actually specifies `ENF_NOPUSH` followed by the operation.

After executing the filter command list, a non-zero value (true) left on top of the stack (or an empty stack) causes the incoming packet to be accepted and a zero value (false) causes the packet to be rejected. (If the filter exits as the result of a short-circuit operator, the top-of-stack value is ignored.) Specifying an undefined operation or action in the command list or performing an illegal operation or action (such as pushing a shortword offset past the end of the packet or executing a binary operator with fewer than two shortwords on the stack) causes a filter to reject the packet.

pfmod(7M)

EXAMPLES

The packet filter module is not dependent on any particular device driver or module but is commonly used with datalink drivers such as the Ethernet driver. If the underlying datalink driver supports the Data Link Provider Interface (DLPI) message set, the appropriate STREAMS DLPI messages must be issued to attach the stream to a particular hardware device and bind a datalink address to the stream before the underlying driver will route received packets upstream. Refer to the DLPI Version 2 specification for details on this interface.

The reverse ARP daemon program may use code similar to the following fragment to construct a filter that rejects all but RARP packets. That is, it accepts only packets whose Ethernet type field has the value `ETHERTYPE_REVARP`.

```
struct ether_header eh;          /* used only for offset values */
struct packetfilt pf;
register ushort_t *fwp = pf.Pf_Filter;
ushort_t offset;
int fd;
/*
 * Push packet filter streams module.
 */
if (ioctl(fd, I_PUSH, "pfmod") < 0)
    syserr("pfmod");

/*
 * Set up filter.  Offset is the displacement of the Ethernet
 * type field from the beginning of the packet in units of
 * ushort_ts.
 */
offset = ((uint_t) &eh.ether_type - (uint_t) &eh.ether_dhost) /
    sizeof (us_short);
*fwp++ = ENF_PUSHPWORD + offset;
*fwp++ = ENF_PUSHLIT;
*fwp++ = htons(ETHERTYPE_REVARP);
*fwp++ = ENF_EQ;
pf.Pf_FilterLen = fwp - &pf.Pf_Filter[0];
```

This filter can be abbreviated by taking advantage of the ability to combine actions and operations:

```
*fwp++ = ENF_PUSHPWORD + offset;
*fwp++ = ENF_PUSHLIT | ENF_EQ;
*fwp++ = htons(ETHERTYPE_REVARP);
```

SEE ALSO [bufmod\(7M\)](#), [dlpi\(7P\)](#), [pullupmsg\(9F\)](#)

NAME	pipemod – STREAMS pipe flushing module
DESCRIPTION	<p>The typical stream is composed of a stream head connected to modules and terminated by a driver. Some stream configurations such as pipes and FIFOs do not have a driver and hence certain features commonly supported by the driver need to be provided by other means. Flushing is one such feature, and it is provided by the pipemod module.</p> <p>Pipes and FIFOs in their simplest configurations only have stream heads. A write side is connected to a read side. This remains true when modules are pushed. The twist occurs at a point known as the mid-point. When an M_FLUSH message is passed from a write queue to a read queue the FLUSHR and/or FLUSHW bits have to be switched. The mid-point of a pipe is not always easily detectable, especially if there are numerous modules pushed on either end of the pipe. In that case there needs to be a mechanism to intercept all message passing through the stream. If the message is an M_FLUSH message and it is at the mid-point, the flush bits need to be switched. This bit switching is handled by the pipemod module.</p> <p>pipemod should be pushed onto a pipe or FIFO where flushing of any kind will take place. The pipemod module can be pushed on either end of the pipe. The only requirement is that it is pushed onto an end that previously did not have modules on it. That is, pipemod must be the first module pushed onto a pipe so that it is at the mid-point of the pipe itself.</p> <p>The pipemod module handles only M_FLUSH messages. All other messages are passed on to the next module using the putnext () utility routine. If an M_FLUSH message is passed to pipemod and the FLUSHR and FLUSHW bits are set, the message is not processed but is passed to the next module using the putnext () routine. If only the FLUSHR bit is set, the FLUSHR bit is turned off and the FLUSHW bit is set. The message is then passed on to the next module using putnext (). Similarly, if the FLUSHW bit is the only bit set in the M_FLUSH message, the FLUSHW bit is turned off and the FLUSHR bit is turned on. The message is then passed to the next module on the stream.</p> <p>The pipemod module can be pushed on any stream that desires the bit switching. It must be pushed onto a pipe or FIFO if any form of flushing must take place.</p>
SEE ALSO	<i>STREAMS Programming Guide</i>

pm(7D)

NAME	pm – Power Management driver						
SYNOPSIS	/dev/pm						
DESCRIPTION	<p>The Power Management (pm) driver provides an interface for applications to configure devices within the system for Power Management. The interface is provided through ioctl(2) commands. The pm driver may be accessed using /dev/pm.</p>						
Power Management Framework	<p>The Power Management framework model allows the system to be viewed as a collection of devices. Each device is a collection of components that comprise the smallest power manageable units. The device driver controls the definition of a device’s power manageable components.</p> <p>A component can either be <i>busy</i> or <i>idle</i> at the current power level. Normally, the Power Management framework takes an <i>idle</i> component to the next lower power level. The Power Management framework uses two factors to determine this transition: the component must have been idle for at least the threshold time, and the device to which the component belongs must satisfy any dependency requirements. A dependency occurs when a device requires another device to be power managed before it can be power managed. Dependencies occur on a per device basis: when a dependency exists, no components of a device may be managed unless all the devices it depends upon are first power managed.</p> <p>Using the commands below, an application may take control of the Power Management of a device from the Power Management framework driver and manage the transition of device power levels directly.</p> <p>For this set of ioctl commands, <i>arg</i> (see ioctl(2)) points to a structure of type pm_req defined in <sys/pm.h>:</p> <pre>typedef struct pm_req { char *physpath; /* physical path of device to configure */ /* see libdevinfo(3LIB) */ int component; /* the component of the device */ int value; /* power level, threshold value, or count */ void *data; /* command-dependent variable sized data */ size_t datasize; /* size of data buffer */ } pm_req_t;</pre> <p>The fields should contain the following data:</p> <table><tr><td><i>physpath</i></td><td>Pointer to the physical path of a device. See libdevinfo(3LIB). For example, for the device /devices/pseudo/pm@0:pm the <i>physpath</i> value would be /pseudo/pm@0.</td></tr><tr><td><i>component</i></td><td>Non-negative integer specifying which component is being configured. The numbering starts at zero.</td></tr><tr><td><i>value</i></td><td>Non-negative integer specifying the threshold value in seconds or the desired power level, or the number of levels being specified.</td></tr></table>	<i>physpath</i>	Pointer to the physical path of a device. See libdevinfo(3LIB). For example, for the device /devices/pseudo/pm@0:pm the <i>physpath</i> value would be /pseudo/pm@0.	<i>component</i>	Non-negative integer specifying which component is being configured. The numbering starts at zero.	<i>value</i>	Non-negative integer specifying the threshold value in seconds or the desired power level, or the number of levels being specified.
<i>physpath</i>	Pointer to the physical path of a device. See libdevinfo(3LIB). For example, for the device /devices/pseudo/pm@0:pm the <i>physpath</i> value would be /pseudo/pm@0.						
<i>component</i>	Non-negative integer specifying which component is being configured. The numbering starts at zero.						
<i>value</i>	Non-negative integer specifying the threshold value in seconds or the desired power level, or the number of levels being specified.						

data Pointer to a buffer which contains or receives variable-sized data, such as the name of a device upon which this device has a dependency.

size Size of the data buffer.

Not all fields are used in each command.

PM_DIRECT_PM

The device named by *physpath* is disabled from being power managed by the framework. The caller will power manage the device directly using the PM_DIRECT_NOTIFY, PM_GET_TIME_IDLE and PM_GET_CURRENT_POWER, PM_GET_FULL_POWER and PM_SET_CURRENT_POWER commands. If the device needs to have its power level changed either because its driver calls pm_raise_power(9F), pm_lower_power(9F), or pm_power_has_changed(9F) or because the device is the parent of another device that is changing power level or a device that this device depends on is changing power level, then the power level change of the device will be blocked and the caller will be notified as described below for the PM_DIRECT_NOTIFY command.

Error codes:

EBUSY Device already disabled for Power Management by framework.

EPERM Caller is neither superuser nor effective group ID of 0.

PM_RELEASE_DIRECT_PM

The device named by *physpath* (which must have been the target of a PM_DIRECT_PM command) is re-enabled for Power Management by the framework.

Error codes:

EINVAL Device component out of range.

PM_DIRECT_NOTIFY PM_DIRECT_NOTIFY_WAIT

These commands allow the process that is directly power managing a device to be notified of events that could change the power level of the device. When such an event occurs, this command returns information about the event.

arg (see ioctl(2)) points to a structure of type pm_state_change defined in <sys/pm.h>:

```
typedef struct pm_state_change {
    char    *physpath;    /* device which has changed state */
    int     component;   /* which component changed state */
#ifdef _BIG_ENDIAN
    ushort_t flags;      /* PSC_EVENT_LOST, PSC_ALL_LOWEST */
    ushort_t event;      /* type of event */
#else
    ushort_t event;      /* type of event */
    ushort_t flags;      /* PSC_EVENT_LOST, PSC_ALL_LOWEST */
#endif
    time_t  timestamp;   /* time of state change */
};
```

pm(7D)

```
int    old_level;    /* power level changing from */
int    new_level;    /* power level changing to */
size_t size;        /* size of buffer physpath points to */
} pm_state_change_t;
```

When an event occurs, the struct pointed to by *arg* is filled in. If the event type is `PSC_PENDING_CHANGE`, then the information in the rest of the struct describes an action that the framework would have taken if the device were not directly power managed by the caller. The caller is responsible for completing the indicated level changes using `PM_SET_CURRENT_POWER` below.

An event type of `PSC_HAS_CHANGED` indicates that the driver for the directly power managed device has called `pm_power_has_changed(9F)` due to the device changing power on its own. It is provided to allow the caller to track the power state of the device.

The system keeps events in a circular buffer. If the buffer overflow, the oldest events are lost and when the event that next follows a lost event is retrieved it will have `PSC_EVENT_LOST` set in flags.

`PM_DIRECT_NOTIFY` returns `EWOULDBLOCK` if no event is pending, and `PM_DIRECT_NOTIFY_WAIT` blocks until an event is available.

`pm` also supports the `poll(2)` interface. When an event is pending a `poll(2)` call that includes a file descriptor for `/dev/pm` and that has `POLLIN` or `POLLRDNORM` set in its event mask will return.

`PM_SET_CURRENT_POWER`

Component *component* of the device named by *physpath* (which must contain the physical path of a device against which the process has issued a `PM_DIRECT_PM` command) is set to power level *value*. If all components of the device named by *physpath* were at level 0, *value* is non-zero and some device has a dependency on this device, then all components of that device will be brought to full power before this command returns. Similarly, if the parent of the target device is powered off, then it will be brought up as needed before this command returns. When `PM_SET_CURRENT_POWER` is issued against a device, the resulting power change is included in the event list for `PM_DIRECT_NOTIFY`.

Error codes:

<code>EINVAL</code>	Device component out of range, or power level < 0.
<code>EIO</code>	Failed to power device or its ancestors or the devices on which this device has dependency or their ancestors. Note that this may not indicate a failure, the device driver may have rejected the command as inappropriate because the component has become busy.
<code>EPERM</code>	Caller has not previously issued a successful <code>PM_DIRECT_PM</code> command against this device.

PM_GET_FULL_POWER

The highest supported power level of component *component* of the device named by *physpath* is returned.

PM_GET_CURRENT_POWER

The current power level of component *component* of the device named by *physpath* is returned.

Error codes:

EAGAIN Device component power level is not currently known.

PM_GET_TIME_IDLE

PM_GET_TIME_IDLE returns the number of seconds that component *component* of the device named by *physpath* has been idle. If the device is not idle, then 0 is returned.

Note that because the state of the device may change between the time the process issues the **PM_GET_TIME_IDLE** command and the time the process issues a **PM_SET_CURRENT_POWER** command to reduce the power level of an idle component, the process must be prepared to deal with a **PM_SET_CURRENT_POWER** command returning failure because the driver has rejected the command as inappropriate because the device component has become busy. This can be differentiated from other types of failures by issuing the **PM_GET_TIME_IDLE** command again to see if the component has become busy.

ERRORS Upon error, the commands will return `-1`, and set *errno*. In addition to the error codes listed above by command, the following error codes are common to all commands:

EFAULT Bad address passed in as argument.

ENODEV Device is not power manageable, or device is not configured.

ENXIO Too many opens attempted.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface stability	Unstable

SEE ALSO `pmconfig(1M)`, `intro(2)`, `ioctl(2)`, `libdevinfo(3LIB)`, `power.conf(4)`, `attributes(5)`, `attach(9E)`, `detach(9E)`, `power(9E)`, `pm_busy_component(9F)`, `pm_idle_component(9F)`, `pm_lower_power(9F)`, `pm_power_has_changed(9F)`, `pm_raise_power(9F)`

Writing Device Drivers

poll(7d)

NAME	<code>poll</code> – driver for fast poll on many file descriptors
SYNOPSIS	<pre>#include <sys/devpoll.h> int fd = open("/dev/poll", O_RDWR); ssize_t n = write(int fd, struct pollfd buf[], int bufsize); int n = ioctl(int fd, DP_POLL, struct dvpoll* arg); int n = ioctl(int fd, DP_ISPOLLED, struct pollfd* pfd);</pre>
PARAMETERS	<p><i>fd</i> Open file descriptor that refers to the <code>/dev/poll</code> driver.</p> <p><i>path</i> <code>/dev/poll</code></p> <p><i>buf</i> Array of <code>pollfd</code> structures.</p> <p><i>bufsize</i> Size of <i>buf</i> in bytes.</p> <p><i>arg</i> Pointer to <code>pollcall</code> structure.</p> <p><i>pfd</i> Pointer to <code>pollfd</code> structure.</p>
DESCRIPTION	<p>Note – The <code>/dev/poll</code> device, associated driver and corresponding manpages may be removed in a future Solaris release. For similar functionality in the event ports framework, see <code>port_create(3C)</code>.</p> <p>The <code>/dev/poll</code> driver is a special driver that enables you to monitor multiple sets of polled file descriptors. By using the <code>/dev/poll</code> driver, you can efficiently poll large numbers of file descriptors. Access to the <code>/dev/poll</code> driver is provided through <code>open(2)</code>, <code>write(2)</code>, and <code>ioctl(2)</code> system calls.</p> <p>Writing an array of <code>pollfd</code> struct to the <code>/dev/poll</code> driver has the effect of adding these file descriptors to the monitored <code>poll</code> file descriptor set represented by the <i>fd</i>. To monitor multiple file descriptor sets, open the <code>/dev/poll</code> driver multiple times. Each <i>fd</i> corresponds to one set. For each <code>pollfd</code> struct entry (defined in <code>sys/poll.h</code>):</p> <pre>struct pollfd { int fd; short events; short revents; }</pre> <p>The <i>fd</i> field specifies the file descriptor being polled. The <i>events</i> field indicates the interested <code>poll</code> events on the file descriptor. If a <code>pollfd</code> array contains multiple <code>pollfd</code> entries with the same <i>fd</i> field, the "events" field in each <code>pollfd</code> entry is OR'ed. A special <code>POLLREMOVE</code> event in the <i>events</i> field of the <code>pollfd</code> structure removes the <i>fd</i> from the monitored set. The <i>revents</i> field is not used. Write returns the number of bytes written successfully or -1 when write fails.</p> <p>The <code>DP_POLL</code> <code>ioctl</code> is used to retrieve returned <code>poll</code> events occurred on the polled file descriptors in the monitored set represented by <i>fd</i>. <i>arg</i> is a pointer to the <code>dvpoll</code> structures which are defined as follows:</p> <pre>struct dvpoll { struct pollfd* dp_fds; int dp_nfds;</pre>

```

    int dp_timeout;
}

```

The `dp_fds` points to a buffer that holds an array of returned `pollfd` structures. The `dp_nfds` field specifies the size of the buffer in terms of the number of `pollfd` entries it contains. The `dp_nfds` field also indicates the maximum number of file descriptors from which poll information can be obtained. If there is no interested events on any of the polled file descriptors, the `DP_POLL` ioctl call will wait `dp_timeout` milliseconds before returning. If `dp_timeout` is 0, the ioctl call returns immediately. If `dp_timeout` is -1, the call blocks until an interested poll events is available or the call is interrupted. Upon return, if the ioctl call has failed, -1 is returned. The memory content pointed by `dp_fds` is not modified. A return value 0 means the ioctl is timed out. In this case, the memory content pointed by `dp_fds` is not modified. If the call is successful, it returns the number of valid `pollfd` entries in the array pointed by `dp_fds`; the contents of the rest of the buffer is undefined. For each valid `pollfd` entry, the `fd` field indicates the file descriptor on which the polled events happened. The `events` field is the user specified poll events. The `revents` field contains the events occurred. -1 is returned if the call fails.

`DP_ISPOLLED` ioctl allows you to query if a file descriptor is already in the monitored set represented by `fd`. The `fd` field of the `pollfd` structure indicates the file descriptor of interest. The `DP_ISPOLLED` ioctl returns 1 if the file descriptor is in the set. The `events` field contains 0. The `revents` field contains the currently polled events. The ioctl returns 0 if the file descriptor is not in the set. The `pollfd` structure pointed by `pdfd` is not modified. The ioctl returns a -1 if the call fails.

EXAMPLES

The following example shows how `/dev/poll` may be used.

```

{
    ...
    /*
     * open the driver
     */
    if ((wfd = open("/dev/poll", O_RDWR)) < 0) {
        exit(-1);
    }
    pollfd = (struct pollfd* )malloc(sizeof(struct pollfd) * MAXBUF);
    if (pollfd == NULL) {
        close(wfd);
        exit(-1);
    }
    /*
     * initialize buffer
     */
    for (i = 0; i < MAXBUF; i++) {
        pollfd[i].fd = fds[i];
        pollfd[i].events = POLLIN;
        pollfd[i].revents = 0;
    }
    if (write(wfd, &pollfd[0], sizeof(struct pollfd) * MAXBUF) !=
        sizeof(struct pollfd) * MAXBUF) {
        perror("failed to write all pollfds");
        close (wfd);
        free(pollfd);
    }
}

```

poll(7d)

```
        exit(-1);
    }
    /*
     * read from the devpoll driver
     */
    dopoll.dp_timeout = -1;
    dopoll.dp_nfds = MAXBUF;
    dopoll.dp_fds = pollfd;
    result = ioctl(wfd, DP_POLL, &dopoll);
    if (result < 0) {
        perror("/dev/poll ioctl DP_POLL failed");
        close (wfd);
        free(pollfd);
        exit(-1);
    }
    for (i = 0; i < result; i++) {
        read(dopoll.dp_fds[i].fd, rbuf, STRLEN);
    }
    ...
}
```

The following example is part of a test program which shows how `DP_ISPOLLED()` `ioctl` may be used.

```
{
    ...

    loopcnt = 0;
    while (loopcnt < ITERATION) {
        rn = random();
        rn %= RANGE;
        if (write(fds[rn], TESTSTRING, strlen(TESTSTRING)) !=
            strlen(TESTSTRING)) {
            perror("write to fifo failed.");
            close (wfd);
            free(pollfd);
            error = 1;
            goto out1;
        }
        dpfd.fd = fds[rn];
        dpfd.events = 0;
        dpfd.revents = 0;
        result = ioctl(wfd, DP_ISPOLLED, &dpfd);
        if (result < 0) {
            perror("/dev/poll ioctl DP_ISPOLLED failed");
            printf("errno = %d\n", errno);
            close (wfd);
            free(pollfd);
            error = 1;
            goto out1;
        }
        if (result != 1) {
            printf("DP_ISPOLLED returned incorrect result: %d.\n",
                result);
            close (wfd);
            free(pollfd);
            error = 1;
        }
    }
}
```

```

        goto out1;
    }
    if (dpfd.fd != fds[rn]) {
        printf("DP_ISPOLLED returned wrong fd %d, expect %d\n",
            dpfd.fd, fds[rn]);
        close (wfd);
        free(pollfd);
        error = 1;
        goto out1;
    }

    if (dpfd.revents != POLLIN) {
        printf("DP_ISPOLLED returned unexpected revents %d\n",
            dpfd.revents);
        close (wfd);
        free(pollfd);
        error = 1;
        goto out1;
    }
    if (read(dpfd.fd, rbuf, strlen(TESTSTRING)) !=
        strlen(TESTSTRING)) {
        perror("read from fifo failed");
        close (wfd);
        free(pollfd);
        error = 1;
        goto out1;
    }
    loopcnt++;
}

```

ERRORS	EACCES	A process does not have permission to access the content cached in /dev/poll.
	EINTR	A signal was caught during the execution of the <code>ioctl(2)</code> function.
	EFAULT	The request argument requires a data transfer to or from a buffer pointed to by <i>arg</i> , but <i>arg</i> points to an illegal address.
	EINVAL	The request or <i>arg</i> parameter is not valid for this device, or field of the <code>dvpoll</code> struct pointed by <i>arg</i> is not valid (for example, <code>dp_nfds</code> is greater than <code>{OPEN_MAX}</code>).
	ENXIO	The <code>O_NONBLOCK</code> flag is set, the named file is a FIFO, the <code>O_WRONLY</code> flag is set, and no process has the file open for reading; or the named file is a character special or block special file and the device associated with this special file does not exist.

ATTRIBUTES See `attributes(5)` for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86

poll(7d)

Availability	SUNWcar (Solaris) SUNWcsr, SUNWcsu (Solaris on x86) SUNWhea (header files)
Interface Stability	Obsolete
MT-Level	Safe

SEE ALSO open(2), poll(2), write(2), attributes(5)

NOTES The `/dev/poll` API is particularly beneficial to applications that poll a large number of file descriptors repeatedly. Applications will exhibit the best performance gain if the polled file descriptor list rarely change.

When using the `/dev/poll` driver, you should remove a closed file descriptor from a monitored poll set. Failure to do so may result in a `POLLNVAL` `revents` being returned for the closed file descriptor. When a file descriptor is closed but not removed from the monitored set, and is reused in subsequent open of a different device, you will be polling the device associated with the reused file descriptor. In a multithreaded application, careful coordination among threads doing close and `DP_POLL` ioctl is recommended for consistent results.

The `/dev/poll` driver caches a list of polled file descriptors, which are specific to a process. Therefore, the `/dev/poll` file descriptor of a process will be inherited by its child process, just like any other file descriptors. But the child process will have very limited access through this inherited `/dev/poll` file descriptor. Any attempt to write or do ioctl by the child process will result in an `EACCES` error. The child process should close the inherited `/dev/poll` file descriptor and open its own if desired.

The `/dev/poll` driver does not yet support polling. Polling on a `/dev/poll` file descriptor will result in `POLLERR` being returned in the `revents` field of `pollfd` structure.

NAME	prnio – generic printer interface
SYNOPSIS	<code>#include <sys/prnio.h></code>
DESCRIPTION	<p>The <code>prnio</code> generic printer interface defines <code>ioctl</code> commands and data structures for printer device drivers.</p> <p><code>prnio</code> defines and provides facilities for five basic phases of the printing process:</p> <ul style="list-style-type: none"> ■ Identification — Retrieve device information/attributes ■ Setup — Set device attributes ■ Transfer — Transfer data to or from the device ■ Cleanup — Transfer phase conclusion ■ Abort — Transfer phase interruption <p>During the Identification phase, the application retrieves a set of device capabilities and additional information using the <code>PRNIOC_GET_IFCAP</code>, <code>PRNIOC_GET_STATUS</code>, <code>PRNIOC_GET_TIMEOUTS</code>, <code>PRNIOC_GET_IFINFO</code> and <code>PRNIOC_GET_1284_DEVID</code> commands.</p> <p>During the Setup phase the application sets some interface attributes and probably resets the printer as described in the <code>PRNIOC_SET_IFCAP</code>, <code>PRNIOC_SET_TIMEOUTS</code> and <code>PRNIOC_RESET</code> sections.</p> <p>During the Transfer phase, data is transferred in a forward (host to peripheral) or reverse direction (peripheral to host). Transfer is accomplished using <code>write(2)</code> and <code>read(2)</code> system calls. For <code>prnio</code> compliant printer drivers, forward transfer support is mandatory, while reverse transfer support is optional. Applications can also use <code>PRNIOC_GET_STATUS</code> and <code>PRNIOC_GET_1284_STATUS</code> commands during the transfer to monitor the device state.</p> <p>The Cleanup phase is accomplished by closing the device using <code>close(2)</code>. Device drivers supporting <code>prnio</code> may set non-zero error code as appropriate. Applications should explicitly <code>close(2)</code> a device before exiting and check <code>errno</code> value.</p> <p>The Abort phase is accomplished by interrupting the <code>write(2)</code> and <code>read(2)</code> system calls. The application can perform some additional cleanup during the Abort phase as described in <code>PRNIOC_GET_IFCAP</code> section.</p>
IOCTLS	<p><code>PRNIOC_GET_IFCAP</code></p> <p>Application can retrieve printer interface capabilities using this command. The <code>ioctl(2)</code> argument is a pointer to <code>uint_t</code>, a bit field representing a set of properties and services provided by a printer driver. Set bit means supported capability. The following values are defined:</p> <p><code>PRN_BIDI</code> - When this bit is set, the interface operates in a bidirectional mode, instead of forward-only mode.</p> <p><code>PRN_HOTPLUG</code> - If this bit is set, the interface allows device hot-plugging.</p> <p><code>PRN_1284_DEVID</code> - If this bit is set, the device is capable of returning 1284 device ID (see <code>PRNIOC_GET_1284_DEVID</code>.)</p>

prnio(7I)

PRN_1284_STATUS - If this bit is set, the device driver can return device status lines (see `PRNIOC_GET_1284_STATUS`). Some devices support this ioctl in unidirectional mode only.

PRN_TIMEOUTS - If this bit is set the peripheral may stall during the transfer phase and the driver can timeout and return from the `write(2)` and `read(2)` returning the number of bytes that have been transferred. If `PRN_TIMEOUTS` is set, the driver supports this functionality and the timeout values can be retrieved and modified via the `PRNIOC_GET_TIMEOUTS` and `PRNIOC_SET_TIMEOUTS` ioctls. Otherwise, applications can implement their own timeouts and abort phase.

PRN_STREAMS - This bit impacts the application abort phase behaviour. If the device claimed `PRN_STREAMS` capability, the application must issue an `I_FLUSH` ioctl(2) before `close(2)` to dismiss the untransferred data. Only `STREAMS` drivers can support this capability.

PRNIOC_SET_IFCAP

This ioctl can be used to change interface capabilities. The argument is a pointer to `uint_t` bit field that is described in detail in the `PRNIOC_GET_IFCAP` section. Capabilities should be set one at a time; otherwise the command will return `EINVAL`. The following capabilities can be changed by this ioctl:

PRN_BIDI - When this capability is set, the interface operates in a bidirectional mode, instead of forward-only mode. Devices that support only one mode will not return error; applications should use `PRNIOC_GET_IFCAP` to check if the mode was successfully changed. Because some capabilities may be altered as a side effect of changing other capabilities, this command should be followed by `PRNIOC_GET_IFCAP`.

PRNIOC_GET_IFINFO

This command can be used to retrieve printer interface info string, which is an arbitrary format string usually describing the bus type. The argument is a pointer to `struct prn_interface_info` as described below.

```
struct prn_interface_info {
    uint_t    if_len;    /* length of buffer */
    uint_t    if_rlen;  /* actual info length */
    char      *if_data; /* buffer address */
};
```

The application allocates a buffer and sets `if_data` and `if_len` values to its address and length, respectively. The driver returns the string to this buffer and sets `if_len` to its length. If `if_len` is less than `if_rlen`, the driver must return the first `if_len` bytes of the string. The application may then repeat the command with a bigger buffer.

Although `prnio` does not limit the contents of the interface info string, some values are recommended and defined in `<sys/prnio.h>` by the following macros:

`PRN_PARALLEL` - Centronics or *IEEE 1284* compatible devices
`PRN_SERIAL` - EIA-232/EIA-485 serial ports

PRN_USB - Universal Serial Bus printers

PRN_1394 - *IEEE 1394* peripherals

Printer interface info string is for information only: no implications should be made from its value.

PRNIOC_RESET

Some applications may want to reset the printer state during Setup and/or Cleanup phase using PRNIOC_RESET command. Reset semantics are device-specific, and in general, applications using this command should be aware of the printer type.

Each prnio compliant driver is required to accept this request, although performed actions are completely driver-dependent. More information on the PRNIOC_RESET implementation for the particular driver is available in the corresponding man page and printer manual.

PRNIOC_GET_1284_DEVID

This command can be used to retrieve printer device ID as defined by *IEEE 1284-1994*. The `ioctl(2)` argument is a pointer to `struct prn_1284_device_id` as described below.

```
struct prn_1284_device_id {
    uint_t      id_len;    /* length of buffer */
    uint_t      id_rlen;   /* actual ID length */
    char        *id_data;  /* buffer address */
};
```

For convenience, the two-byte length field is not considered part of device ID string and is not returned in the user buffer. Instead, `id_rlen` value shall be set to $(\text{length} - 2)$ by the driver, where length is the ID length field value. If buffer length is less than `id_rlen`, the driver returns the first `id_len` bytes of the ID.

The printer driver must return the most up-to-date value of the device ID.

PRNIOC_GET_STATUS

This command can be used by applications to retrieve current device status. The argument is a pointer to `uint_t`, where the status word is returned. Status is a combination of the following bits:

PRN_ONLINE - For devices that support PRN_HOTPLUG capability, this bit is set when the device is online, otherwise the device is offline. Devices without PRN_HOTPLUG support should always have this bit set.

PRN_READY - This bit indicates if the device is ready to receive/send data.

Applications may use this bit for an outbound flow control

PRNIOC_GET_1284_STATUS

Devices that support PRN_1284_STATUS capability accept this `ioctl` to retrieve the device status lines defined in *IEEE 1284* for use in Compatibility mode. The following bits may be set by the driver:

PRN_1284_NOFAULT - Device is not in error state

prnio(7I)

PRN_1284_SELECT - Device is selected
PRN_1284_PE - Paper error
PRN_1284_BUSY - Device is busy

PRNIOC_GET_TIMEOUTS

This command retrieves current transfer timeout values for the driver. The argument is a pointer to `struct prn_timeouts` as described below.

```
struct prn_timeouts {
    uint_t    tmo_forward; /* forward transfer timeout */
    uint_t    tmo_reverse; /* reverse transfer timeout */
};
```

`tmo_forward` and `tmo_reverse` define forward and reverse transfer timeouts in seconds. This command is only valid for drivers that support `PRN_TIMEOUTS` capability.

PRNIOC_SET_TIMEOUTS

This command sets current transfer timeout values for the driver. The argument is a pointer to `struct prn_timeouts`. See `PRNIOC_GET_TIMEOUTS` for description of this structure. This command is only valid for drivers that support `PRN_TIMEOUTS` capability.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, IA
Interface Stability	Evolving

SEE ALSO `close(2)`, `ioctl(2)`, `read(2)`, `write(2)`, `attributes(5)`, `ecpp(7D)`, `usbprn(7D)`, `lp(7D)`

IEEE Std 1284-1994

NAME	profile – DTrace profile interrupt provider
DESCRIPTION	<p>The <code>profile</code> driver is a DTrace dynamic tracing provider that adds time-based interrupt event sources that can be used as DTrace probes.</p> <p>Each profile event source is a time-based interrupt firing every fixed, specified time interval. You can use these probes to sample some aspect of system state every unit time and the samples can then be used to infer system behavior. If the sampling rate is high, or the sampling time is long, an accurate inference is possible. By using the DTrace facility to bind arbitrary actions to probes, you can use the <code>profile</code> provider to sample practically anything in the system. For example, you could sample the state of the current thread, the CPU state, or the current machine instruction each time a probe fires.</p> <p>The <code>profile</code> driver is not a public interface and you access the instrumentation offered by this provider through DTrace. Refer to the <i>Solaris Dynamic Tracing Guide</i> for a description of the public documented interfaces available for the DTrace facility and the probes offered by the profile provider.</p>
ATTRIBUTES	See <code>attributes(5)</code> for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdtrp
Interface Stability	Private

SEE ALSO `dtrace(1M)`, `attributes(5)`, `dtrace(7D)`
Solaris Dynamic Tracing Guide

ptem(7M)

NAME	ptem – STREAMS Pseudo Terminal Emulation module
SYNOPSIS	<pre>int ioctl(fd, I_PUSH, "ptem");</pre>
DESCRIPTION	<p>ptem is a STREAMS module that, when used in conjunction with a line discipline and pseudo terminal driver, emulates a terminal.</p> <p>The ptem module must be pushed (see <code>I_PUSH</code>, <code>streamio(7I)</code>) onto the slave side of a pseudo terminal STREAM, before the <code>ldterm(7M)</code> module is pushed.</p> <p>On the write-side, the <code>TCSETA</code>, <code>TCSETAF</code>, <code>TCSETAW</code>, <code>TCGETA</code>, <code>TCSETS</code>, <code>TCSETSW</code>, <code>TCSETSF</code>, <code>TCGETS</code>, <code>TCSBRK</code>, <code>JWINSIZE</code>, <code>TIOCGWINSZ</code>, and <code>TIOCSWINSZ</code> termio <code>ioctl(2)</code> messages are processed and acknowledged. If remote mode is not in effect, ptem handles the <code>TIOCSTI</code> <code>ioctl</code> by copying the argument bytes into an <code>M_DATA</code> message and passing it back up the read side. Regardless of the remote mode setting, ptem acknowledges the <code>ioctl</code> and passes a copy of it downstream for possible further processing. A hang up (that is, <code>stty 0</code>) is converted to a zero length <code>M_DATA</code> message and passed downstream. Termio <code>cflags</code> and window row and column information are stored locally one per stream. <code>M_DELAY</code> messages are discarded. All other messages are passed downstream unmodified.</p> <p>On the read-side all messages are passed upstream unmodified with the following exceptions. All <code>M_READ</code> and <code>M_DELAY</code> messages are freed in both directions. A <code>TCSBRK</code> <code>ioctl</code> is converted to an <code>M_BREAK</code> message and passed upstream and an acknowledgement is returned downstream. A <code>TIOCSIGNAL</code> <code>ioctl</code> is converted into an <code>M_PCSIG</code> message, and passed upstream and an acknowledgement is returned downstream. Finally a <code>TIOCREMOTE</code> <code>ioctl</code> is converted into an <code>M_CTL</code> message, acknowledged, and passed upstream; the resulting mode is retained for use in subsequent <code>TIOCSTI</code> parsing.</p>
FILES	<sys/ptem.h>
SEE ALSO	<code>stty(1)</code> , <code>ioctl(2)</code> , <code>ldterm(7M)</code> , <code>pckt(7M)</code> , <code>streamio(7I)</code> , <code>termio(7I)</code> <i>STREAMS Programming Guide</i>

NAME	ptm – STREAMS pseudo-tty master driver
DESCRIPTION	<p>The pseudo-tty subsystem simulates a terminal connection, where the master side represents the terminal and the slave represents the user process's special device end point. In order to use the pseudo-tty subsystem, a node for the master side driver <code>/dev/ptmx</code> and <code>N</code> number of nodes for the slave driver must be installed. See pts(7D). The master device is set up as a cloned device where its major device number is the major for the clone device and its minor device number is the major for the <code>ptm</code> driver. There are no nodes in the file system for master devices. The master pseudo driver is opened using the <code>open(2)</code> system call with <code>/dev/ptmx</code> as the device parameter. The clone open finds the next available minor device for the <code>ptm</code> major device.</p> <p>A master device is available only if it and its corresponding slave device are not already open. When the master device is opened, the corresponding slave device is automatically locked out. Only one open is allowed on a master device. Multiple opens are allowed on the slave device. After both the master and slave have been opened, the user has two file descriptors which are the end points of a full duplex connection composed of two streams which are automatically connected at the master and slave drivers. The user may then push modules onto either side of the stream pair.</p> <p>The master and slave drivers pass all messages to their adjacent queues. Only the <code>M_FLUSH</code> needs some processing. Because the read queue of one side is connected to the write queue of the other, the <code>FLUSHR</code> flag is changed to the <code>FLUSHW</code> flag and vice versa. When the master device is closed an <code>M_HANGUP</code> message is sent to the slave device which will render the device unusable. The process on the slave side gets the <code>errno EIO</code> when attempting to write on that stream but it will be able to read any data remaining on the stream head read queue. When all the data has been read, <code>read()</code> returns 0 indicating that the stream can no longer be used. On the last close of the slave device, a 0-length message is sent to the master device. When the application on the master side issues a <code>read()</code> or <code>getmsg()</code> and 0 is returned, the user of the master device decides whether to issue a <code>close()</code> that dismantles the pseudo-terminal subsystem. If the master device is not closed, the pseudo-tty subsystem will be available to another user to open the slave device.</p> <p>If <code>O_NONBLOCK</code> or <code>O_NDELAY</code> is set, read on the master side returns -1 with <code>errno</code> set to <code>EAGAIN</code> if no data is available, and write returns -1 with <code>errno</code> set to <code>EAGAIN</code> if there is internal flow control.</p>
IOCTLS	The master driver supports the <code>ISPTM</code> and <code>UNLKPT</code> ioctls that are used by the functions <code>grantpt(3C)</code> , <code>unlockpt(3C)</code> and <code>ptsname(3C)</code> . The ioctl <code>ISPTM</code> determines whether the file descriptor is that of an open master device. On success, it returns the 0. The ioctl <code>UNLKPT</code> unlocks the master and slave devices. It returns 0 on success. On failure, the <code>errno</code> is set to <code>EINVAL</code> indicating that the master device is not open.
FILES	<code>/dev/ptmx</code> master clone device

ptm(7D)

/dev/pts/M slave devices (M = 0 -> N-1)

SEE ALSO [grantpt\(3C\)](#), [ptsname\(3C\)](#), [unlockpt\(3C\)](#), [pckt\(7M\)](#), [pts\(7D\)](#)

STREAMS Programming Guide

NAME	pts – STREAMS pseudo-tty slave driver
DESCRIPTION	<p>The pseudo-tty subsystem simulates a terminal connection, where the master side represents the terminal and the slave represents the user process's special device end point. In order to use the pseudo-tty subsystem, a node for the master side driver <code>/dev/ptmx</code> and <code>N</code> nodes for the slave driver (<code>N</code> is determined at installation time) must be installed. The names of the slave devices are <code>/dev/pts/M</code> where <code>M</code> has the values 0 through <code>N-1</code>. When the master device is opened, the corresponding slave device is automatically locked out. No user may open that slave device until its permissions are adjusted and the device unlocked by calling functions <code>grantpt(3C)</code> and <code>unlockpt(3C)</code>. The user can then invoke the open system call with the name that is returned by the <code>ptsname(3C)</code> function. See the example below.</p> <p>Only one open is allowed on a master device. Multiple opens are allowed on the slave device. After both the master and slave have been opened, the user has two file descriptors which are end points of a full duplex connection composed of two streams automatically connected at the master and slave drivers. The user may then push modules onto either side of the stream pair. The user needs to push the <code>psem(7M)</code> and <code>ldterm(7M)</code> modules onto the slave side of the pseudo-terminal subsystem to get terminal semantics.</p> <p>The master and slave drivers pass all messages to their adjacent queues. Only the <code>M_FLUSH</code> needs some processing. Because the read queue of one side is connected to the write queue of the other, the <code>FLUSHR</code> flag is changed to the <code>FLUSHW</code> flag and vice versa. When the master device is closed an <code>M_HANGUP</code> message is sent to the slave device which will render the device unusable. The process on the slave side gets the error <code>EIO</code> when attempting to write on that stream but it will be able to read any data remaining on the stream head read queue. When all the data has been read, <code>read</code> returns 0 indicating that the stream can no longer be used. On the last close of the slave device, a 0-length message is sent to the master device. When the application on the master side issues a <code>read()</code> or <code>getmsg()</code> and 0 is returned, the user of the master device decides whether to issue a <code>close()</code> that dismantles the pseudo-terminal subsystem. If the master device is not closed, the pseudo-tty subsystem will be available to another user to open the slave device. Since 0-length messages are used to indicate that the process on the slave side has closed and should be interpreted that way by the process on the master side, applications on the slave side should not write 0-length messages. If that occurs, the write returns 0, and the 0-length message is discarded by the <code>psem</code> module.</p> <p>The standard STREAMS system calls can access the pseudo-tty devices. The slave devices support the <code>O_NDELAY</code> and <code>O_NONBLOCK</code> flags.</p>
EXAMPLES	<pre>int fdm fds; char *slavename; extern char *ptsname(); fdm = open("/dev/ptmx", O_RDWR); /* open master */ grantpt(fdm); /* change permission of slave */ unlockpt(fdm); /* unlock slave */ slavename = ptsname(fdm); /* get name of slave */</pre>

pts(7D)

```
fds = open(slavename, O_RDWR); /* open slave */
ioctl(fds, I_PUSH, "ptem"); /* push ptem */
ioctl(fds, I_PUSH, "ldterm"); /* push ldterm*/
```

FILES /dev/ptmx master clone device
/dev/pts/M slave devices (M = 0 -> N-1)

SEE ALSO [grantpt\(3C\)](#), [ptsname\(3C\)](#), [unlockpt\(3C\)](#), [ldterm\(7M\)](#), [ptm\(7D\)](#), [ptem\(7M\)](#)

STREAMS Programming Guide

NAME	pty – pseudo-terminal driver
DESCRIPTION	<p>The <code>pty</code> driver provides support for a pair of devices collectively known as a <i>pseudo-terminal</i>. The two devices comprising a pseudo-terminal are known as a <i>controller</i> and a <i>slave</i>. The slave device distinguishes between the <code>B0</code> baud rate and other baud rates specified in the <code>c_cflag</code> word of the <code>termios</code> structure, and the <code>CLOCAL</code> flag in that word. It does not support any of the other <code>termio(7I)</code> device control functions specified by flags in the <code>c_cflag</code> word of the <code>termios</code> structure and by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code>, or <code>INPCK</code> flags in the <code>c_iflag</code> word of the <code>termios</code> structure, as these functions apply only to asynchronous serial ports. All other <code>termio(7I)</code> functions must be performed by STREAMS modules pushed atop the driver; when a slave device is opened, the <code>ldterm(7M)</code> and <code>ttcompat(7M)</code> STREAMS modules are automatically pushed on top of the stream, providing the standard <code>termio(7I)</code> interface.</p> <p>Instead of having a hardware interface and associated hardware that supports the terminal functions, the functions are implemented by another process manipulating the controller device of the pseudo-terminal.</p> <p>The controller and the slave devices of the pseudo-terminal are tightly connected. Any data written on the controller device is given to the slave device as input, as though it had been received from a hardware interface. Any data written on the slave terminal can be read from the controller device (rather than being transmitted from a UAR).</p> <p>By default, 48 pseudo-terminal pairs are configured as follows:</p> <pre>/dev/pty[p-r][0-9a-f] controller devices /dev/tty[p-r][0-9a-f] slave devices</pre>
IOCTLS	<p>The standard set of <code>termio</code> <code>ioctl</code>s are supported by the slave device. None of the bits in the <code>c_cflag</code> word have any effect on the pseudo-terminal, except that if the baud rate is set to <code>B0</code>, it will appear to the process on the controller device as if the last process on the slave device had closed the line; thus, setting the baud rate to <code>B0</code> has the effect of “hanging up” the pseudo-terminal, just as it has the effect of “hanging up” a real terminal.</p> <p>There is no notion of “parity” on a pseudo-terminal, so none of the flags in the <code>c_iflag</code> word that control the processing of parity errors have any effect. Similarly, there is no notion of a “break”, so none of the flags that control the processing of breaks, and none of the <code>ioctl</code>s that generate breaks, have any effect.</p> <p>Input flow control is automatically performed; a process that attempts to write to the controller device will be blocked if too much unconsumed data is buffered on the slave device. The input flow control provided by the <code>IXOFF</code> flag in the <code>c_iflag</code> word is not supported.</p> <p>The delays specified in the <code>c_oflag</code> word are not supported.</p> <p>As there are no modems involved in a pseudo-terminal, the <code>ioctl</code>s that return or alter the state of modem control lines are silently ignored.</p>

pty(7D)

A few special `ioctl`s are provided on the controller devices of pseudo-terminals to provide the functionality needed by applications programs to emulate real hardware interfaces:

TIOCSTOP	The argument is ignored. Output to the pseudo-terminal is suspended, as if a <code>STOP</code> character had been typed.												
TIOCSTART	The argument is ignored. Output to the pseudo-terminal is restarted, as if a <code>START</code> character had been typed.												
TIOCPKT	<p>The argument is a pointer to an <code>int</code>. If the value of the <code>int</code> is non-zero, <i>packet</i> mode is enabled; if the value of the <code>int</code> is zero, packet mode is disabled. When a pseudo-terminal is in packet mode, each subsequent <code>read(2)</code> from the controller device will return data written on the slave device preceded by a zero byte (symbolically defined as <code>TIOCPKT_DATA</code>), or a single byte reflecting control status information. In the latter case, the byte is an inclusive-or of zero or more of the bits:</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="vertical-align: top;"><code>TIOCPKT_FLUSHREAD</code></td> <td>whenever the read queue for the terminal is flushed.</td> </tr> <tr> <td style="vertical-align: top;"><code>TIOCPKT_FLUSHWRITE</code></td> <td>whenever the write queue for the terminal is flushed.</td> </tr> <tr> <td style="vertical-align: top;"><code>TIOCPKT_STOP</code></td> <td>whenever output to the terminal is stopped using <code>^S</code>.</td> </tr> <tr> <td style="vertical-align: top;"><code>TIOCPKT_START</code></td> <td>whenever output to the terminal is restarted.</td> </tr> <tr> <td style="vertical-align: top;"><code>TIOCPKT_DOSTOP</code></td> <td>whenever <code>XON/XOFF</code> flow control is enabled after being disabled; it is considered "enabled" when the <code>IXON</code> flag in the <code>c_iflag</code> word is set, the <code>VSTOP</code> member of the <code>c_cc</code> array is <code>^S</code> and the <code>VSTART</code> member of the <code>c_cc</code> array is <code>^Q</code>.</td> </tr> <tr> <td style="vertical-align: top;"><code>TIOCPKT_NOSTOP</code></td> <td>whenever <code>XON/XOFF</code> flow control is disabled after being enabled.</td> </tr> </table>	<code>TIOCPKT_FLUSHREAD</code>	whenever the read queue for the terminal is flushed.	<code>TIOCPKT_FLUSHWRITE</code>	whenever the write queue for the terminal is flushed.	<code>TIOCPKT_STOP</code>	whenever output to the terminal is stopped using <code>^S</code> .	<code>TIOCPKT_START</code>	whenever output to the terminal is restarted.	<code>TIOCPKT_DOSTOP</code>	whenever <code>XON/XOFF</code> flow control is enabled after being disabled; it is considered "enabled" when the <code>IXON</code> flag in the <code>c_iflag</code> word is set, the <code>VSTOP</code> member of the <code>c_cc</code> array is <code>^S</code> and the <code>VSTART</code> member of the <code>c_cc</code> array is <code>^Q</code> .	<code>TIOCPKT_NOSTOP</code>	whenever <code>XON/XOFF</code> flow control is disabled after being enabled.
<code>TIOCPKT_FLUSHREAD</code>	whenever the read queue for the terminal is flushed.												
<code>TIOCPKT_FLUSHWRITE</code>	whenever the write queue for the terminal is flushed.												
<code>TIOCPKT_STOP</code>	whenever output to the terminal is stopped using <code>^S</code> .												
<code>TIOCPKT_START</code>	whenever output to the terminal is restarted.												
<code>TIOCPKT_DOSTOP</code>	whenever <code>XON/XOFF</code> flow control is enabled after being disabled; it is considered "enabled" when the <code>IXON</code> flag in the <code>c_iflag</code> word is set, the <code>VSTOP</code> member of the <code>c_cc</code> array is <code>^S</code> and the <code>VSTART</code> member of the <code>c_cc</code> array is <code>^Q</code> .												
<code>TIOCPKT_NOSTOP</code>	whenever <code>XON/XOFF</code> flow control is disabled after being enabled.												
TIOCREMOTE	The argument is a pointer to an <code>int</code> . If the value of the <code>int</code> is non-zero, <i>remote</i> mode is enabled; if the value of the <code>int</code> is zero, remote mode is disabled. This mode can be enabled or disabled independently of packet mode. When a pseudo-terminal is in remote mode, input to the slave device of the pseudo-terminal is flow controlled and not input edited (regardless of the mode the slave side of the pseudo-terminal). Each write to the controller device produces a record boundary for the process reading the slave device. In normal usage, a write of data is like the data typed as a line on the terminal; a write of 0 bytes is like typing an EOF												

character. Note: this means that a process writing to a pseudo-terminal controller in *remote* mode must keep track of line boundaries, and write only one line at a time to the controller. If, for example, it were to buffer up several NEWLINE characters and write them to the controller with one `write()`, it would appear to a process reading from the slave as if a single line containing several NEWLINE characters had been typed (as if, for example, a user had typed the LNEXT character before typing all but the last of those NEWLINE characters). Remote mode can be used when doing remote line editing in a window manager, or whenever flow controlled input is required.

EXAMPLES

```
#include <fcntl.h>
#include <sys/termios.h>

int fdm fds;
fdm = open("/dev/ptyp0, O_RDWR); /* open master */
fds = open("/dev/ttyp0, O_RDWR); /* open slave */
```

FILES /dev/pty [p-z] [0-9a-f] pseudo-terminal controller devices
 /dev/tty [p-z] [0-9a-f] pseudo-terminal slave devices

SEE ALSO rlogin(1), rlogind(1M), ldterm(7M), termio(7I), ttcompat(7M),

NOTES It is apparently not possible to send an EOT by writing zero bytes in TIOCREMOTE mode.

qfe(7d)

NAME	qfe – SUNW,qfe Quad Fast-Ethernet device driver
SYNOPSIS	/dev/qfe
DESCRIPTION	<p>The SUNW,qfe Quad Fast-Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, dlpi(7P), over a SUNW,qfe Quad Fast-Ethernet controller. Multiple SUNW,qfe controllers installed within the system are supported by the driver. The qfe driver provides basic support for the SUNW,qfe hardware. It is used to handle the SUNW,qfe device. Functions include chip initialization, frame transit and receive, multicast and promiscuous support, and error recovery and reporting.</p>
SUNW,qfe	<p>The SUNW,qfe device provides a 100Base-TX networking interface. There are two types of SUNW,qfe device: one supporting Sbus and the other supporting the PCI bus interface. The Sbus SUNW,qfe device uses Sun's FEPS ASIC, which provides the Sbus interface and MAC functions. The PCI SUNW,qfe device uses Sun's PFEX ASIC to provide the PCI interface and MAC functions. Both connect with the 100Base-TX on-board transceiver, which connects to a RJ45 connector to provide the Physical layer functions and external connection.</p> <p>The 100Base-TX standard specifies an "auto-negotiation" protocol to automatically select the mode and speed of operation. The internal transceiver is capable of doing auto-negotiation with the remote-end of the link (link partner) and receives the capabilities of the remote end. It selects the Highest Common Denominator mode of operation based on the priorities. It also supports forced-mode of operation where the driver can select the mode of operation.</p>
APPLICATION PROGRAMMING INTERFACE qfe and DLPI	<p>The cloning character-special device /dev/qfe is used to access all SUNW,qfe controllers installed within the system.</p> <p>The qfe driver is a "style 2" data link service provider. All M_PROTO and M_PCPROTO type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <sys/dlpi.h>. Refer to dlpi(7P) for more information. An explicit DL_ATTACH_REQ message by the user is required to associate the opened stream with a particular device (ppa). The ppa ID is interpreted as an unsigned long data type and indicates the corresponding device instance (unit) number. The driver returns an error (DL_ERROR_ACK) if the ppa field value does not correspond to a valid device instance number for this system. The device is initialized on first attach and de-initialized (stopped) at last detach.</p> <p>The values returned by the driver in the DL_INFO_ACK primitive in response to the DL_INFO_REQ from the user are as follows:</p> <ul style="list-style-type: none">■ The maximum SDU is 1500 (ETHERMTU - defined in <sys/ethernet.h>).■ The minimum SDU is 0.■ The dlsap address length is 8.■ The MAC type is DL_ETHER.

- The `sap` length values is `-2` meaning the physical address component is followed immediately by a 2 byte `sap` component within the DLSAP address.
- The service mode is `DL_CLDLS`.
- No optional quality of service (QOS) support is included at present so the QOS fields are 0.
- The provider style is `DL_STYLE2`.
- The version is `DL_VERSION_2`.
- The broadcast address value is Ethernet/IEEE broadcast address (`0xFFFFFFFF`).

Once in the `DL_ATTACHED` state, the user must send a `DL_BIND_REQ` to associate a particular *service access pointer* SAP with the stream. The `qfe` driver interprets the `sap` field within the `DL_BIND_REQ` as an Ethernet “type” therefore valid values for the `sap` field are in the `[0-0xFFFF]` range. Only one Ethernet type can be bound to the stream at any time.

If the user selects a `sap` with a value of 0, the receiver will be in “802.3 mode”. All frames received from the media having a “type” field in the range `[0-1500]` are assumed to be 802.3 frames and are routed up all open streams which are bound to `sap` value 0. If more than one stream is in “802.3 mode” then the frame will be duplicated and routed up multiple streams as `DL_UNITDATA_IND` messages.

In transmission, the driver checks the `sap` field of the `DL_BIND_REQ` if the `sap` value is 0, and if the destination type field is in the range `[0-1500]`. If either is true, the driver computes the length of the message, not including initial `M_PROTO` mblk (message block), of all subsequent `DL_UNITDATA_REQ` messages and transmits 802.3 frames that have this value in the MAC frame header length field.

The `qfe` driver DLSAP address format consists of the 6 byte physical (Ethernet) address component followed immediately by the 2 byte `sap` (type) component producing an 8 byte DLSAP address. Applications should *not* hardcode to this particular implementation-specific DLSAP address format but use information returned in the `DL_INFO_ACK` primitive to compose and decompose DLSAP addresses. The `sap` length, full DLSAP length, and `sap`/physical ordering are included within the `DL_INFO_ACK`. The physical address length can be computed by subtracting the `sap` length from the full DLSAP address length or by issuing the `DL_PHYS_ADDR_REQ` to obtain the current physical address associated with the stream.

Once in the `DL_BOUND` state, the user may transmit frames on the Ethernet by sending `DL_UNITDATA_REQ` messages to the `qfe` driver. The `qfe` driver will route received Ethernet frames up all those open and bound streams having a `sap` which matches the Ethernet type as `DL_UNITDATA_IND` messages. Received Ethernet frames are duplicated and routed up multiple open streams if necessary. The DLSAP address contained within the `DL_UNITDATA_REQ` and `DL_UNITDATA_IND` messages consists of both the `sap` (type) and physical (Ethernet) components.

qfe(7d)

<p>qfe Primitives</p>	<p>In addition to the mandatory connectionless DLPI message set the driver also supports the following primitives.</p> <p>The <code>DL_ENABMULTI_REQ</code> and <code>DL_DISABMULTI_REQ</code> primitives enable or disable reception of individual multicast group addresses. A set of multicast addresses may be iteratively created and modified on a per-stream basis using these primitives. The driver accepts these primitives in any state following <code>DL_ATTACHED</code>.</p> <p>The <code>DL_PROMISCON_REQ</code> and <code>DL_PROMISCOFF_REQ</code> primitives with the <code>DL_PROMISC_PHYS</code> flag set in the <code>dl_level</code> field enables or disables reception of all frames on the media ("promiscuous mode"), including frames generated by the local host.</p> <p>When used with the <code>DL_PROMISC_SAP</code> flag set this enables or disables reception of all <code>sap</code> (Ethernet type) values. When used with the <code>DL_PROMISC_MULTI</code> flag set this enables or disables reception of all multicast group addresses. The effect of each is always on a per-stream basis and independent of the other <code>sap</code> and physical level configurations on this stream or other streams.</p> <p>The <code>DL_PHYS_ADDR_REQ</code> primitive returns the 6 octet Ethernet address currently associated (attached) to the stream in the <code>DL_PHYS_ADDR_ACK</code> primitive. This primitive is valid only in states following a successful <code>DL_ATTACH_REQ</code>.</p> <p>The <code>DL_SET_PHYS_ADDR_REQ</code> primitive changes the 6 octet Ethernet address currently associated (attached) to this stream. The credentials of the process which originally opened this stream must be root. Otherwise <code>EPERM</code> is returned in the <code>DL_ERROR_ACK</code>. This primitive is destructive in that it affects all other current and future streams attached to this device. An <code>M_ERROR</code> is sent up all other streams attached to this device when this primitive is successful on this stream. Once changed, all streams subsequently opened and attached to this device will obtain this new physical address. Once changed, the physical address will remain until this primitive is used to change the physical address again or the system is rebooted, whichever comes first.</p>
<p>qfe Driver</p>	<p>By default, the <code>qfe</code> driver performs "auto-negotiation" to select the mode and speed of the link.</p> <p>The link can be in one of the four following modes:</p> <ul style="list-style-type: none">■ 100 Mbps, full-duplex■ 100 Mbps, half-duplex■ 10 Mbps, full-duplex■ 10 Mbps, half-duplex <p>These speeds and modes are described in the 100Base-TX standard.</p> <p>The auto-negotiation protocol automatically selects:</p> <ul style="list-style-type: none">■ Operation mode (half-duplex or full-duplex)■ Speed (100 Mbps or 10 Mbps)

The auto-negotiation protocol does the following:

- Gets all the modes of operation supported by the Link Partner
- Advertises its capabilities to the Link Partner
- Selects the highest common denominator mode of operation based on the priorities.
- The highest priority is given to the 100 Mbps, full-duplex; lowest priority is given to 10 Mbps, half-duplex.

The *100Base-TX transceiver* is capable of all of the operating speeds and modes listed above. By default, auto-negotiation is used to select the speed and the mode of the link and the common mode of operation with the link partner.

Sometimes, the user may want to select the speed and mode of the link. The SUNW,qfe device supports programmable "IPG" (Inter-Packet Gap) parameters `ipg1` and `ipg2`. By default, the driver sets `ipg1` to 8 `byte-times` and `ipg2` to 4 `byte-times` (which are the standard values). Sometimes, the user may want to alter these values depending on whether the driver supports 10 Mbps or 100 Mbps and accordingly, IPG will be set to 9.6 or 0.96 microseconds.

qfe Parameter List

The qfe driver provides for setting and getting various parameters for the SUNW,qfe device. The parameter list includes:

- current transceiver status
- current link status
- inter-packet gap
- local transceiver capabilities
- link partner capabilities

The local transceiver has two sets of capabilities: one set reflects the capabilities of the hardware, which are read-only (RO) parameters, and the second set, which reflects the values chosen by the user, is used in speed selection. There are read/write (RW) capabilities. At boot time, these two sets of capabilities will be the same. The Link Partner capabilities are also read-only parameters because the current default value of these parameters can only be read and cannot be modified.

FILES `/dev/qfe` qfe special character device
 `/kernel/drv/qfe.conf` system wide default device driver properties

SEE ALSO `ndd(1M)`, `netstat(1M)`, `driver.conf(4)`, `dlpi(7P)`

qlc(7D)

NAME	qlc – ISP2200 and ISP2300 Family Fibre Channel host bus adapter driver						
SYNOPSIS	SUNW, qlc						
DESCRIPTION	<p>The qlc host bus adapter driver is a Sun Fibre Channel transport layer-compliant nexus driver for the Qlogic ISP2200, ISP2200A, ISP2310 and ISP2312 adapters. These adapters support Fibre Channel SCSI and IP Protocols, FC-AL public loop profile, point-to-point fabric connection and Fibre Channel service classes two and three.</p> <p>The qlc driver interfaces with the Sun Fibre Channel transport layer to support the standard functions provided by the SCSA interface. It supports auto request sense and tagged queueing by default. The driver requires that all devices have unique hard addresses in private loop configurations. Devices with conflicting hard addresses are not accessible.</p>						
FILES	<p>/kernel/drv/qlc 32-bit ELF kernel module (x86)</p> <p>/kernel/drv/amd64/qlc 64-bit ELF kernel module (x86)</p> <p>/kernel/drv/sparcv9/qlc 64-bit ELF kernel module (SPARC)</p> <p>/kernel/drv/qlc.conf Driver configuration file</p>						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>SPARC, x86</td></tr><tr><td>Availability</td><td>SUNWqlc</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC, x86	Availability	SUNWqlc
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Architecture	SPARC, x86						
Availability	SUNWqlc						
SEE ALSO	<p>prtconf(1M), driver.conf(4), fcp(7D), fp(7D)</p> <p><i>Writing Device Drivers</i></p> <p><i>ANSI X3.230:1994, Fibre Channel Physical Signaling (FC-PH)</i></p> <p><i>Project 1134-D, Fibre Channel Generic Services (FC-GS-2)</i></p> <p><i>ANSI X3.269-1996, Fibre Channel Arbitrated Loop (FC-AL)</i></p> <p><i>ANSI X3.270-1996, Fibre Channel Protocol for SCSI (FCP-SCSI)</i></p> <p><i>ANSI X3.270-1996, SCSI-3 Architecture Model (SAM)</i></p> <p><i>Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)</i></p> <p><i>Fabric Loop Attachment (FC-FLA)</i></p>						

qlc(7D)

ISP2200 Firmware Interface Specification, QLogic Corporation

2300 Series Firmware Specification, QLogic Corporation

quotactl(7I)

NAME	quotactl – manipulate disk quotas														
SYNOPSIS	<pre>#include <sys/fs/ufs_quota.h> int ioctl(int fd, Q_QUOTACTL, struct quotctl *qp)</pre>														
DESCRIPTION	<p>This <code>ioctl()</code> call manipulates disk quotas. <i>fd</i> is the file descriptor returned by the <code>open()</code> system call after opening the <code>quotas</code> file (located in the root directory of the filesystem running quotas.) <code>Q_QUOTACTL</code> is defined in <code>/usr/include/sys/fs/ufs_quota.h</code>. <i>qp</i> is the address of the <code>quotctl</code> structure which is defined as</p> <pre>struct quotctl { int op; uid_t uid; caddr_t addr; };</pre> <p><i>op</i> indicates an operation to be applied to the user ID <i>uid</i>. (See below.) <i>addr</i> is the address of an optional, command specific, data structure which is copied in or out of the system. The interpretation of <i>addr</i> is given with each value of <i>op</i> below.</p> <table><tr><td><code>Q_QUOTAON</code></td><td>Turn on quotas for a file system. <i>addr</i> points to the full pathname of the <code>quotas</code> file. <i>uid</i> is ignored. It is recommended that <i>uid</i> have the value of 0. This call is restricted to the super-user.</td></tr><tr><td><code>Q_QUOTAOFF</code></td><td>Turn off quotas for a file system. <i>addr</i> and <i>uid</i> are ignored. It is recommended that <i>addr</i> have the value of <code>NULL</code> and <i>uid</i> have the value of 0. This call is restricted to the super-user.</td></tr><tr><td><code>Q_GETQUOTA</code></td><td>Get disk quota limits and current usage for user <i>uid</i>. <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code><sys/fs/ufs_quota.h></code>). Only the super-user may get the quotas of a user other than himself.</td></tr><tr><td><code>Q_SETQUOTA</code></td><td>Set disk quota limits and current usage for user <i>uid</i>. <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code>sys/fs/ufs_quota.h</code>). This call is restricted to the super-user.</td></tr><tr><td><code>Q_SETQLIM</code></td><td>Set disk quota limits for user <i>uid</i>. <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code>sys/fs/ufs_quota.h</code>). This call is restricted to the super-user.</td></tr><tr><td><code>Q_SYNC</code></td><td>Update the on-disk copy of quota usages for this file system. <i>addr</i> and <i>uid</i> are ignored.</td></tr><tr><td><code>Q_ALLSYNC</code></td><td>Update the on-disk copy of quota usages for all file systems with active quotas. <i>addr</i> and <i>uid</i> are ignored.</td></tr></table>	<code>Q_QUOTAON</code>	Turn on quotas for a file system. <i>addr</i> points to the full pathname of the <code>quotas</code> file. <i>uid</i> is ignored. It is recommended that <i>uid</i> have the value of 0. This call is restricted to the super-user.	<code>Q_QUOTAOFF</code>	Turn off quotas for a file system. <i>addr</i> and <i>uid</i> are ignored. It is recommended that <i>addr</i> have the value of <code>NULL</code> and <i>uid</i> have the value of 0. This call is restricted to the super-user.	<code>Q_GETQUOTA</code>	Get disk quota limits and current usage for user <i>uid</i> . <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code><sys/fs/ufs_quota.h></code>). Only the super-user may get the quotas of a user other than himself.	<code>Q_SETQUOTA</code>	Set disk quota limits and current usage for user <i>uid</i> . <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code>sys/fs/ufs_quota.h</code>). This call is restricted to the super-user.	<code>Q_SETQLIM</code>	Set disk quota limits for user <i>uid</i> . <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code>sys/fs/ufs_quota.h</code>). This call is restricted to the super-user.	<code>Q_SYNC</code>	Update the on-disk copy of quota usages for this file system. <i>addr</i> and <i>uid</i> are ignored.	<code>Q_ALLSYNC</code>	Update the on-disk copy of quota usages for all file systems with active quotas. <i>addr</i> and <i>uid</i> are ignored.
<code>Q_QUOTAON</code>	Turn on quotas for a file system. <i>addr</i> points to the full pathname of the <code>quotas</code> file. <i>uid</i> is ignored. It is recommended that <i>uid</i> have the value of 0. This call is restricted to the super-user.														
<code>Q_QUOTAOFF</code>	Turn off quotas for a file system. <i>addr</i> and <i>uid</i> are ignored. It is recommended that <i>addr</i> have the value of <code>NULL</code> and <i>uid</i> have the value of 0. This call is restricted to the super-user.														
<code>Q_GETQUOTA</code>	Get disk quota limits and current usage for user <i>uid</i> . <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code><sys/fs/ufs_quota.h></code>). Only the super-user may get the quotas of a user other than himself.														
<code>Q_SETQUOTA</code>	Set disk quota limits and current usage for user <i>uid</i> . <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code>sys/fs/ufs_quota.h</code>). This call is restricted to the super-user.														
<code>Q_SETQLIM</code>	Set disk quota limits for user <i>uid</i> . <i>addr</i> is a pointer to a <code>dqblk</code> structure (defined in <code>sys/fs/ufs_quota.h</code>). This call is restricted to the super-user.														
<code>Q_SYNC</code>	Update the on-disk copy of quota usages for this file system. <i>addr</i> and <i>uid</i> are ignored.														
<code>Q_ALLSYNC</code>	Update the on-disk copy of quota usages for all file systems with active quotas. <i>addr</i> and <i>uid</i> are ignored.														
RETURN VALUES	<p>This <code>ioctl()</code> returns:</p> <table><tr><td>0</td><td>on success.</td></tr><tr><td>-1</td><td>on failure and sets <code>errno</code> to indicate the error.</td></tr></table>	0	on success.	-1	on failure and sets <code>errno</code> to indicate the error.										
0	on success.														
-1	on failure and sets <code>errno</code> to indicate the error.														

ERRORS	EFAULT	<i>addr</i> is invalid.	
	EINVAL	The kernel has not been compiled with the QUOTA option. <i>op</i> is invalid.	
	ENOENT	The quotas file specified by <i>addr</i> does not exist.	
	EPERM	The call is privileged and the calling process did not assert {PRIV_SYS_MOUNT} in the effective set.	
	ESRCH	No disk quota is found for the indicated user. Quotas have not been turned on for this file system.	
	EUSERS	The quota table is full.	
	If <i>op</i> is Q_QUOTAON, <code>ioctl()</code> may set <code>errno</code> to:		
	EACCES	The quota file pointed to by <i>addr</i> exists but is not a regular file. The quota file pointed to by <i>addr</i> exists but is not on the file system pointed to by <i>special</i> .	
EIO	Internal I/O error while attempting to read the quotas file pointed to by <i>addr</i> .		
FILES	<code>/usr/include/sys/fs/ufs_quota.h</code>	quota-related structure/function definitions and defines	
SEE ALSO	quota(1M), quotacheck(1M), quotaon(1M), <code>getrlimit(2)</code> , <code>mount(2)</code>		
BUGS	There should be some way to integrate this call with the resource limit interface provided by <code>setrlimit()</code> and <code>getrlimit(2)</code> .		
	This call is incompatible with Melbourne quotas.		

qus(7D)

NAME	qus – Qlogic Ultra3 SCSI ISP10160 Host Bus Adapter Driver
SYNOPSIS	<pre>pci@pci-slot/scsi@4 - Scsi bus-1 pci@pci-slot/scsi@5 - Scsi bus-2</pre>
DESCRIPTION	<p>The ISP10160 host bus adapter is a SCSI-compliant nexus driver that supports Qlogic ISP10160 SCSI chips on the PCI bus. The ISP10160 is an intelligent SCSI host bus adapter chip that reduces the amount of CPU overhead used in a SCSI transfer.</p> <p>The qus driver supports the standard functions provided by the SCSI interface. The driver supports tagged and untagged queuing, fast, fast-20, fast-40, fast-80, and auto request sense, but does not support linked commands.</p>
CONFIGURATION	<p>You configure the qus driver by defining properties in <code>qus.conf</code> which override the global SCSI settings. Supported properties are <code>scsi-options</code>, <code>target<n>-scsi-options</code>, <code>scsi-reset-delay</code>, <code>scsi-watchdog-tick</code>, <code>scsi-tag-age-limit</code>, <code>scsi-initiator-id</code>, and <code>scsi-selection-timeout</code>.</p> <p><code>target<n>-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target<n></code>. <code><n></code> is a hex value that can vary from 0 to f. Refer to <code>scsi_hba_attach(9F)</code> for details.</p>
EXAMPLES	<p>Example 1: SCSI Options:</p> <p>Create a file called <code>/kernel/drv/qus.conf</code>, then add the following line to disable tagged queuing, fast SCSI, and Wide mode for all qus instances:</p> <pre>scsi-options=0x78;</pre> <p>To disable an option for a specific ISP10160 (see <code>driver.conf(4)</code>), do the following:</p> <pre>name="qus" parent="/pci@1f,2000/pci@1" unit-address="4" scsi-options=0x178 target3-scsi-options=0x58 scsi-initiator-id=6;</pre> <p>Note that the default initiator ID in OBP is 7 and that the change to ID 6 occurs at attach time. You may prefer to change the initiator ID in OBP.</p> <p>Example 1 sets <code>scsi-options</code> for target 3 to 0x58 and all other targets on this SCSI bus to 0x178.</p> <p>To determine the physical pathname of the parent, use the <code>/devices</code> tree or follow the link of the logical device name:</p> <pre>example# ls -l /dev/rdisk/c2t0d0s0 lrwxrwxrwx 1 root root 76 Aug 22 13:29 /dev/rdisk/c2t0d0 s0 -> ../../devices/pci@1f,2000/pci@1/scsi@5/sd@0,0:a,raw</pre> <p>To determine the hardware property values, use the output of <code>prtconf(1M)</code> with the <code>-v</code> option:</p>

```

pci, instance #0
  Driver properties:
    name='device_type' type=string items=1 dev=none
    value='pci'
  Hardware properties:
    name='ranges' type=int items=8
    value=82000000.00000000.00100000.82000000.00000000.00100000./
    00000000.00100000
    name='latency-timer' type=int items=1
    value=00000040
    name='cache-line-size' type=int items=1
    value=00000010
scsi, instance #0
  Driver properties:
    name='scsi-selection-timeout' type=int items=1 dev=(249,0)
    value=000000fa
    name='scsi-options' type=int items=1 dev=(249,0)
    value=00107ff8
    name='scsi-watchdog-tick' type=int items=1 dev=(249,0)
    value=0000000a
    name='scsi-tag-age-limit' type=int items=1 dev=(249,0)
    value=00000002
    name='scsi-reset-delay' type=int items=1 dev=(249,0)
    value=00000bb8
  Hardware properties:
    name='cache-line-size' type=int items=1
    value=00000010
  sd (driver not attached)
  st (driver not attached)
scsi, instance #1
  Driver properties:
    name='scsi-selection-timeout' type=int items=1 dev=(249,0)
    value=000000fa
    name='scsi-options' type=int items=1 dev=(249,0)
    value=00107ff8
    name='scsi-watchdog-tick' type=int items=1 dev=(249,0)
    value=0000000a
    name='scsi-tag-age-limit' type=int items=1 dev=(249,0)
    value=00000002
    name='scsi-reset-delay' type=int items=1 dev=(249,0)
    value=00000bb8
  Hardware properties:
    name='cache-line-size' type=int items=1
    value=00000010
  sd (driver not attached)
  st (driver not attached)

```

Example 2: ISP10160 Properties

qus(7D)

The qus driver exports properties indicating (per target) the negotiated transfer speed (target<n>-sync-speed), whether tagged queuing has been enabled (target<n>-TQ), and whether the wide data transfer has been negotiated (target<n>-wide). The sync-speed property value is the data transfer rate in KB/sec. The target-TQ and target-wide properties have no value. The existence of these properties indicate that tagged queuing or wide transfer is enabled. Refer to prtconf(1M) (verbose option) for information on qus properties.

```
scsi, instance #1
  Driver properties:
    name='target2-wide' type=boolean dev=none
    name='target2-TQ' type=boolean dev=none
    name='target2-sync-speed' type=int items=1 dev=none
      value=00027100
    name='target0-wide' type=boolean dev=none
    name='target0-TQ' type=boolean dev=none
    name='target0-sync-speed' type=int items=1 dev=none
      value=00027100
```

To determine the physical pathname of the parent, use the /devices tree or follow the link of the logical device name.

To set scsi-options more specifically per device type, add the following line to the /kernel/drv/qus.conf file:

```
device-type-scsi-options-list =
  "SEAGATE ST32550W", "seagate-scsi-options" ;
  seagate-scsi-options = 0x58;
```

All devices of this specific disk type have scsi-options set to 0x58.

scsi-options specified per target ID has the highest precedence, followed by scsi-options per device type. Global (for all qus instances) scsi-options per bus has the lowest precedence.

You must reboot your system for the specified scsi-options to take effect.

Example 3: Driver Capabilities

To enable some driver features, the target driver must set capabilities in the qus driver. The target driver can query and modify the following capabilities: synchronous, tagged-qing, wide-xfer, auto-rqsense, qfull-retries, qfull-retry-interval. All other capabilities are query only.

By default, tagged-qing, auto-rqsense, and wide-xfer capabilities are disabled, while disconnect, synchronous, and untagged-qing are enabled. These capabilities can have binary values (0 or 1) only. The default value for both qfull-retries and qfull-retry-interval is 10. The qfull-retries capability is a uchar_t (0 to 255), while qfull-retry-interval is a ushort_t (0 to 65535).

qus(7D)

The target driver must enable tagged-qing and wide-xfer explicitly. The untagged-qing capability is always enabled and its value cannot be modified due to the qus driver's ability to queue commands even when tagged-qing is disabled.

When a conflict occurs between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only whom != 0 is supported in the `scsi_ifsetcap(9F)` call.

Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

FILES /kernel/drv/sparcv9/qus64-bitELF kernel module
/kernel/drv/qus.conf Configuration file

ATTRIBUTES See `attributes(5)` for a description of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO `prtconf(1M)`, `driver.conf(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

Writing Device Drivers

ANSI Small Computer System Interface-2 (SCSI-2)

SCSI Parallel Interface-3 (SPI-3)

QLogic Corporation, ISP1080/1280/10160A/12160A Initiator Firmware Interface Specification

QLogic Corporation, ISP10160A/12160A Technical Manual

QLogic Corporation, ISP1080 Technical Manual

QLogic Corporation, ISP1280 Technical Manual

DIAGNOSTICS In addition to being logged, the following messages may appear on the system console.

The first set of messages shown below may be displayed while the qus driver is first trying to attach. All messages in the first set indicate that the qus driver was unable to attach. Each message is preceded by "qus<number>)", where "<number>" is the instance number of the JASPER Host Bus Adapter.

Device in slave-only slot, unused

The SBus device has been placed in a slave-only slot and will not be accessible; move to non- slave-only SBus slot.

qus(7D)

Device is using a hilevel intr, unused

The device was configured with an interrupt level that cannot be used with the qus driver. Check the device.

Failed to alloc soft state

Driver is unable to allocate space for the internal state structure. Driver did not attach to device; SCSI devices are inaccessible.

Bad soft state

Driver requested an invalid internal state structure. Driver did not attach to device; SCSI devices are inaccessible.

Can't reload firmware: failing attach/resume

Driver is unable to reload firmware; check for bad hardware. Driver did not attach to device; SCSI devices are inaccessible.

Can't reset interface during attach/resume

Driver is unable to reset the hardware. Driver did not attach to device; SCSI devices are inaccessible.

Cannot find PCI device-id

Driver is unable to find PCI device-id. Driver did not attach to device; SCSI devices are inaccessible.

Unable to support ISP chip

Driver is unable to support the ISP chip, which is in the interface. Driver did not attach to device; SCSI devices are inaccessible.

Unable to map pci config registers

Driver is unable to map device registers. Check for bad hardware. Driver did not attach to device. SCSI devices are inaccessible.

Unable to attach: check for hardware problem

Driver is unable to attach to the hardware. Driver did not attach to device; SCSI devices are inaccessible.

The following messages can be displayed at any time and are printed with the full device pathname followed by the shorter form described above.

Firmware (<actual-size>) should be <allowed_size> bytes

Firmware id(<id>) verification failed

Firmware length too short

Firmware checksum incorrect

Can't find firmware rev. string

reset/init ISP chip failed

reset/init ISP chip failed

Load ram failed

Start firmware mailbox command failed

Can't set clock rate

Can't get RAM info

These messages indicate a firmware download failure and possible corruption of the firmware. Check the ISP driver.

Chip reset timeout

ISP chip failed to reset in the time allocated. Potential hardware problem.

Bad request pkt payload

The ISP Firmware rejected the packet, indicating that the packet was set up incorrectly. As a result, the `qus` driver will call the target completion routine with the reason of `CMD_TRAN_ERR` set in the `scsi_pkt`. To correctly set up the packet, check the target driver.

Bad request pkt header

The ISP Firmware rejected the packet because the packet was set up incorrectly. As a result, the `qus` driver will call the target completion routine with the reason of `CMD_TRAN_ERR` set in the `scsi_pkt`. To correctly set up the packet, check the target driver.

Target synch. rate reduced. tgt <target-id>**Target<target-id> reducing transfer rate**

These messages indicate that the target is reducing its transfer rate. Reboot the system to obtain the maximum transfer rate.

Failed to Get Features**Chip reset detected**

These messages indicate a possible ISP chip failure. Driver attempts to recover from this condition by reloading and restarting the firmware.

Interface going offline

Although all driver recovery procedure are completed, the interface did not come online and may need replacement.

SCSI Cable/Connection problem

The SCSI cable is faulty or is connected improperly.

Hardware/Firmware error

The ISP chip encountered a firmware error that is probably due to a faulty SCSI cable or improper cable connection. As a result, the `qus` driver attempts to do error recovery by resetting the chip.

Received unexpected SCSI Reset

The ISP chip received an unexpected SCSI Reset and has initiated its own internal error recovery, which returns `scsi_pkt` with reason set to `CMD_RESET`.

Fatal error, resetting interface

The `qus` driver is performing error recovery. As a result, all outstanding commands that have been transported to the `qus` driver will be completed via the `scsi_pkt` completion routine in the target driver with reason of `CMD_RESET` and status of `STAT_BUS_RESET` set in `scsi_pkt`.

Phase skipped: Command completed with good status — Potential data path failure:

Possible SCSI cable failure. Driver may still be able to communicate with the target and continue sending commands. Leaving the system in the same state triggers this message once per hour. System I/O throughput may be reduced. The SCSI may need to be replaced.

qus(7D)

LVD Error detected
SCSI Cable/Connection problem
Bus Not Terminated
f/w initiated BDR fails

These messages indicate a possible cable failure. Outstanding commands are returned with `CMD_RESET` or `CMD_TRAN_ERR` set in `scsi_pkt`.

NAME	ramdisk – RAM disk device driver
SYNOPSIS	<code>ramdisk@0:diskname</code>
DESCRIPTION	The ramdisk driver supports numerous ramdisk devices that are created by the system during the boot process (see <code>boot(1M)</code>) or during normal system operation (see <code>ramdiskadm(1M)</code> for more information).
DEVICE SPECIAL FILES	<p>Each ramdisk can be accessed either as a block device or as a raw device. When accessed as a block device, the normal buffering mechanism is used when reading from and writing to the device, without regard to physical disk records. Accessing the ramdisk as a raw device enables direct transmission between the disk and the read or write buffer. A single read or write call usually results in a single I/O operation, meaning that raw I/O is more efficient when many bytes are transmitted. You can find block files names in <code>/dev/ramdisk</code>. Raw file names are found in <code>/dev/rramdisk</code>.</p> <p>There are no alignment or length restrictions on I/O requests to either block or character devices.</p>
ERRORS	<p>EFAULT The argument features a bad address.</p> <p>EINVAL Invalid argument. EIO. An I/O error occurred.</p> <p>EPERM Cannot create or delete a ramdisk without write permission on <code>/dev/ramdiskctl</code>.</p> <p>ENOTTY The device does not support the requested ioctl function.</p> <p>ENXIO The device did not exist during opening.</p> <p>EBUSY Cannot exclusively open <code>/dev/ramdiskctl</code>. One or more ramdisks are still open.</p> <p>EEXIST A ramdisk with the indicated name already exists.</p> <p>EAGAIN Cannot allocate resource for ramdisk. Try again later.</p>
FILES	<p><code>/dev/ramdisk/diskname</code> Block device for ramdisk named <i>diskname</i>.</p> <p><code>/dev/rramdisk/diskname</code> Raw device for ramdisk name <i>diskname</i></p> <p><code>/kernel/drv/ramdisk</code> 32-bit driver</p> <p><code>/kernel/drv/ramdisk.conf</code> Driver configuration file. (Do not alter).</p> <p><code>/kernel/drv/sparcv9/ramdisk</code> 64-bit driver</p>

ramdisk(7D)

ATTRIBUTES See `attributes(5)` for descriptions of the following attribute:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO `ramdiskadm(1M)`, `fsck(1M)`, `fstyp(1M)`, `mount(1M)`, `newfs(1M)`, `driver.conf(4)`, `filesystem(5)`, `dkio(7I)`

NOTES The percentage of available physical memory that can be allocated to ramdisks is constrained by the variable `rd_percent_physmem`. You can tune the `rd_percent_physmem` variable in `/etc/system`. By default, the percentage of available physical memory that can be allocated to ramdisks is fixed at 25%.

A ramdisk may not be the best possible use of system memory. Accordingly, use ramdisks only when absolutely necessary.

NAME	random, urandom – Strong random number generator device				
SYNOPSIS	<pre>/dev/random /dev/urandom</pre>				
DESCRIPTION	<p>The <code>/dev/random</code> and <code>/dev/urandom</code> files are special files that are a source for random bytes generated by the kernel random number generator device. The <code>/dev/random</code> and <code>/dev/urandom</code> files are suitable for applications requiring high quality random numbers for cryptographic purposes.</p> <p>The generator device produces random numbers from data and devices available to the kernel and estimates the amount of randomness (or "entropy") collected from these sources. The entropy level determines the amount of high quality random numbers that are produced at a given time.</p> <p>Applications retrieve random bytes by reading <code>/dev/random</code> or <code>/dev/urandom</code>. The <code>/dev/random</code> interface returns random bytes only when sufficient amount of entropy has been collected. If there is no entropy to produce the requested number of bytes, <code>/dev/random</code> blocks until more entropy can be obtained. Non-blocking I/O mode can be used to disable the blocking behavior. The <code>/dev/random</code> interface also supports <code>poll(2)</code>. Note that using <code>poll(2)</code> will not increase the speed at which random numbers can be read.</p> <p>Bytes retrieved from <code>/dev/random</code> provide the highest quality random numbers produced by the generator, and can be used to generate long term keys and other high value keying material.</p> <p>The <code>/dev/urandom</code> interface returns bytes regardless of the amount of entropy available. It does not block on a read request due to lack of entropy. While bytes produced by the <code>/dev/urandom</code> interface are of lower quality than bytes produced by <code>/dev/random</code>, they are nonetheless suitable for less demanding and shorter term cryptographic uses such as short term session keys, paddings, and challenge strings.</p> <p>Data can be written to <code>/dev/random</code> and <code>/dev/urandom</code>. Data written to either special file is added to the generator's internal state. Data that is difficult to predict by other users may contribute randomness to the generator state and help improve the quality of future generated random numbers.</p> <p>By default, write access is restricted to the super-user. An administrator may change the default read/write restriction by changing the permissions on the appropriate special files.</p> <p><code>/dev/random</code> collects entropy from providers that are registered with the kernel-level cryptographic framework and implement random number generation routines. The <code>cryptoadm(1M)</code> utility allows an administrator to configure which providers will be used with <code>/dev/random</code>.</p>				
ERRORS	<table border="0"> <tr> <td style="vertical-align: top;">EAGAIN</td> <td>O_NDELAY or O_NONBLOCK was set and no random bytes are available for reading from <code>/dev/random</code>.</td> </tr> <tr> <td style="vertical-align: top;">EINTR</td> <td>A signal was caught while reading and no data was transferred.</td> </tr> </table>	EAGAIN	O_NDELAY or O_NONBLOCK was set and no random bytes are available for reading from <code>/dev/random</code> .	EINTR	A signal was caught while reading and no data was transferred.
EAGAIN	O_NDELAY or O_NONBLOCK was set and no random bytes are available for reading from <code>/dev/random</code> .				
EINTR	A signal was caught while reading and no data was transferred.				

random(7D)

ENOXIO `open(2)` request failed on `/dev/random` because no entropy provider is available.

FILES `/dev/random`
 `/dev/urandom`

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr
Interface Stability	Evolving

SEE ALSO `cryptoadm(1M)`, `open(2)`, `poll(2)`, `attributes(5)`

NOTES `/dev/random` can be configured to use only the hardware-based providers registered with the kernel-level cryptographic framework by disabling the software-based provider using `cryptoadm(1M)`. You can also use `cryptoadm(1M)` to obtain the name of the software-based provider.

Because no entropy is available, disabling all randomness providers causes `read(2)` and `poll(2)` on `/dev/random` to block indefinitely and results in a warning message being logged and displayed on the system console. However, `read(2)` and `poll(2)` on `/dev/random` continue to work in this case.

An implementation of the `/dev/random` and `/dev/urandom` kernel-based random number generator first appeared in Linux 1.3.30.

A `/dev/random` interface for Solaris first appeared as part of the CryptoRand implementation.

NAME	rarp, RARP – Reverse address resolution protocol
DESCRIPTION	<p>You use the RARP protocol to map dynamically between the Internet Protocol (IP) and network interface MAC addresses. RARP is often used to boot a Solaris client. RARP clients include the SPARC boot PROM, x86 boot floppy, SunOS kernel, and <code>ifconfig(1M)</code>. <code>in.rarpd(1M)</code> provides the server-side implementation.</p> <p>RARP request timeout behavior in application-layer clients is governed by the <code>/etc/inet/rarp</code> default file. To tune the number of retries an application attempts before giving up, set the <code>RARP_RETRIES</code> variable in <code>/etc/inet/rarp</code>. If the file is not present or <code>RARP_RETRIES</code> is not initialized within it, applications retry a maximum of five times with a eight second wait between retries.</p>
FILES	<code>/etc/inet/rarp</code>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability (protocol)	Standard
Interface Stability (defaults file)	Unstable
Interface Stability (RARP_RETRIES)	Unstable

SEE ALSO `ifconfig(1M)`, `in.rarpd(1M)`, `arp(7P)`

Reverse Address Resolution Protocol RFC 903. June, 1994 R. Finlayson, T. Mann, J.C. Mogul, M. Theimer

route(7P)

NAME	route – kernel packet forwarding database
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/socket.h> #include <net/if.h> #include <net/route.h> int socket (PF_ROUTE, SOCK_RAW, int protocol) ;</pre>
DESCRIPTION	<p>UNIX provides some packet routing facilities. The kernel maintains a routing information database, which is used in selecting the appropriate network interface when transmitting packets.</p> <p>A user process (or possibly multiple co-operating processes) maintains this database by sending messages over a special kind of socket. This supplants fixed size <code>ioctl(2)</code>'s specified in routing(7P). Routing table changes may only be carried out by the superuser.</p> <p>The operating system may spontaneously emit routing messages in response to external events, such as receipt of a re-direct, or failure to locate a suitable route for a request. The message types are described in greater detail below.</p> <p>Routing database entries come in two flavors: entries for a specific host, or entries for all hosts on a generic subnetwork (as specified by a bit mask and value under the mask). The effect of wildcard or default route may be achieved by using a mask of all zeros, and there may be hierarchical routes.</p> <p>When the system is booted and addresses are assigned to the network interfaces, the internet protocol family installs a routing table entry for each interface when it is ready for traffic. Normally the protocol specifies the route through each interface as a <i>direct</i> connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface is requested to address the packet to the gateway listed in the routing entry, that is, the packet is forwarded.</p> <p>When routing a packet, the kernel attempts to find the most specific route matching the destination. If no entry is found, the destination is declared to be unreachable, and a routing-miss message is generated if there are any listeners on the routing control socket (described below). If there are two different mask and value-under-the-mask pairs that match, the more specific is the one with more bits in the mask. A route to a host is regarded as being supplied with a mask of as many ones as there are bits in the destination.</p> <p>A wildcard routing entry is specified with a zero destination address value, and a mask of all zeroes. Wildcard routes are used when the system fails to find other routes matching the destination. The combination of wildcard routes and routing redirects can provide an economical mechanism for routing traffic.</p> <p>One opens the channel for passing routing control messages by using the socket call shown in the SYNOPSIS section above. There can be more than one routing socket open per system.</p>

Messages are formed by a header followed by a small number of `sockaddrs`, whose length depend on the address family. `sockaddrs` are interpreted by position. An example of a type of message with three addresses might be a CIDR prefix route: Destination, Netmask, and Gateway. The interpretation of which addresses are present is given by a bit mask within the header, and the sequence is least significant to most significant bit within the vector.

Any messages sent to the kernel are returned, and copies are sent to all interested listeners. The kernel provides the process ID of the sender, and the sender may use an additional sequence field to distinguish between outstanding messages. However, message replies may be lost when kernel buffers are exhausted.

The *protocol* parameter specifies which messages an application listening on the routing socket is interested in seeing, based on the the address family of the `sockaddrs` present. Currently, you can specify `AF_INET` and `AF_INET6` to filter the messages seen by the listener, or alternatively, you can specify `AF_UNSPEC` to indicate that the listener is interested in all routing messages.

The kernel may reject certain messages, and will indicate this by filling in the `rtm_errno` field of the `rt_msghdr` struct (see below). The following codes may be returned:

<code>EEXIST</code>	If requested to duplicate an existing entry
<code>ESRCH</code>	If requested to delete a non-existent entry
<code>ENOBUFS</code>	If insufficient resources were available to install a new route.
<code>EPERM</code>	If the calling process does not have appropriate privileges to alter the routing table.

In the current implementation, all routing processes run locally, and the values for `rtm_errno` are available through the normal `errno` mechanism, even if the routing reply message is lost.

A process may avoid the expense of reading replies to its own messages by issuing a `setsockopt(3SOCKET)` call indicating that the `SO_USELOOPBACK` option at the `SOL_SOCKET` level is to be turned off. A process may ignore all messages from the routing socket by doing a `shutdown(3SOCKET)` system call for further input.

If a route is in use when it is deleted, the routing entry is marked down and removed from the routing table, but the resources associated with it are not reclaimed until all references to it are released.

Messages

User processes can obtain information about the routing entry to a specific destination by using a `RTM_GET` message.

Messages include:

```
#define RTM_ADD      0x1    /* Add Route */
#define RTM_DELETE  0x2    /* Delete Route */
#define RTM_CHANGE  0x3    /* Change Metrics, Flags, or Gateway */
```

route(7P)

```
#define RTM_GET      0x4  /* Report Information */
#define RTM_LOSING   0x5  /* Kernel Suspects Partitioning */
#define RTM_REDIRECT 0x6  /* Told to use different route */
#define RTM_MISS     0x7  /* Lookup failed on this address */
#define RTM_LOCK     0x8  /* fix specified metrics */
#define RTM_OLDADD   0x9  /* caused by SIOCADDRT */
#define RTM_OLDDEL   0xa  /* caused by SIOCDELRT */
#define RTM_RESOLVE  0xb  /* request to resolve dst to LL addr */
#define RTM_NEWADDR  0xc  /* address being added to iface */
#define RTM_DELADDR  0xd  /* address being removed from iface */
#define RTM_IFINFO   0xe  /* iface going up/down etc. */
```

A message header consists of:

```
struct rt_msghdr {
    ushort_t rtm_msglen; /* to skip over non-understood messages */
    uchar_t  rtm_version; /* future binary compatibility */
    uchar_t  rtm_type;    /* message type */
    ushort_t rtm_index;  /* index for associated ifp */
    pid_t    rtm_pid;    /* identify sender */
    int      rtm_addrs;  /* bitmask identifying sockaddrs in msg */
    int      rtm_seq;    /* for sender to identify action */
    int      rtm_errno;  /* why failed */
    int      rtm_flags;  /* flags, incl kern & message, e.g., DONE */
    int      rtm_use;    /* from rtenry */
    uint_t   rtm_inits;  /* which values we are initializing */
};
```

```
struct rt_metrics rtm_rmx; /* metrics themselves */
};
```

where

```
struct rt_metrics {
    uint32_t rmx_locks; /* Kernel must leave these values alone */
    uint32_t rmx_mtu;  /* MTU for this path */
    uint32_t rmx_hopcount; /* max hops expected */
    uint32_t rmx_expire; /* lifetime for route, e.g., redirect */
    uint32_t rmx_recvpipe; /* inbound delay-bandwidth product */
    uint32_t rmx_sendpipe; /* outbound delay-bandwidth product */
    uint32_t rmx_ssthresh; /* outbound gateway buffer limit */
    uint32_t rmx_rtt;     /* estimated round trip time */
    uint32_t rmx_rttvar;  /* estimated rtt variance */
    uint32_t rmx_pksent;  /* packets sent using this route */
};
```

/* Flags include the values */

```
#define RTF_UP      0x1  /* route usable */
#define RTF_GATEWAY 0x2  /* destination is a gateway */
#define RTF_HOST    0x4  /* host entry (net otherwise) */
#define RTF_REJECT  0x8  /* host or net unreachable */
#define RTF_DYNAMIC 0x10 /* created dynamically (by redirect) */
#define RTF_MODIFIED 0x20 /* modified dynamically (by redirect) */
#define RTF_DONE    0x40 /* message confirmed */
#define RTF_MASK    0x80 /* subnet mask present */
#define RTF_CLONING 0x100 /* generate new routes on use */
#define RTF_XRESOLVE 0x200 /* external daemon resolves name */
```

```

#define RTF_LLINFO      0x400    /* generated by ARP */
#define RTF_STATIC      0x800    /* manually added */
#define RTF_BLACKHOLE   0x1000   /* just discard pkts (during updates) */
#define RTF_PRIVATE     0x2000   /* do not advertise this route */
#define RTF_PROTO2      0x4000   /* protocol specific routing flag #2 */
#define RTF_PROTO1      0x8000   /* protocol specific routing flag #1 */

/* Specifiers for metric values in rmx_locks and rtm_inits are */

#define RTV_MTU          0x1      /* init or lock _mtu */
#define RTV_HOPCOUNT   0x2      /* init or lock _hopcount */
#define RTV_EXPIRE      0x4      /* init or lock _expire */
#define RTV_RPIPE       0x8      /* init or lock _recvpipe */
#define RTV_SPIPE       0x10     /* init or lock _sendpipe */
#define RTV_SSTHRESH    0x20     /* init or lock _ssthresh */
#define RTV_RTT         0x40     /* init or lock _rtt */
#define RTV_RTTVAR      0x80     /* init or lock _rttvar */

/* Specifiers for which addresses are present in the messages are */

#define RTA_DST          0x1      /* destination sockaddr present */
#define RTA_GATEWAY     0x2      /* gateway sockaddr present */
#define RTA_NETMASK     0x4      /* netmask sockaddr present */
#define RTA_GENMASK     0x8      /* cloning mask sockaddr present */
#define RTA_IFP         0x10     /* interface name sockaddr present */
#define RTA_IFA         0x20     /* interface addr sockaddr present */
#define RTA_AUTHOR      0x40     /* sockaddr for author of redirect */
#define RTA_BRD         0x80     /* for NEWADDR, broadcast or p-p dest addr */

```

SEE ALSO [ioctl\(2\)](#), [setsockopt\(3SOCKET\)](#), [shutdown\(3SOCKET\)](#), [routing\(7P\)](#)

NOTES Some of the metrics may not be implemented and return zero. The implemented metrics are set in `rtm_inits`.

routing(7P)

NAME	routing – system support for packet network routing
DESCRIPTION	<p>The network facilities provide general packet routing. The <code>routing</code> interface described here can be used to maintain the system's IPv4 routing table. It has been maintained for compatibility with older applications. The recommended interface for maintaining the system's routing tables is the routing socket, described at route(7P). The routing socket can be used to manipulate both the IPv4 and IPv6 routing tables of the system. Routing table maintenance may be implemented in applications processes.</p> <p>A simple set of data structures compose a "routing table" used in selecting the appropriate network interface when transmitting packets. This table contains a single entry for each route to a specific network or host. The routing table was designed to support routing for the Internet Protocol (IP), but its implementation is protocol independent and thus it may serve other protocols as well. User programs may manipulate this data base with the aid of two <code>ioctl(2)</code> commands, <code>SIOCADDRT</code> and <code>SIOCDELRT</code>. These commands allow the addition and deletion of a single routing table entry, respectively. Routing table manipulations may only be carried out by privileged user.</p> <p>A routing table entry has the following form, as defined in <code>/usr/include/net/route.h</code>:</p> <pre>struct rtenry { unit_t rt_hash; /* to speed lookups */ struct sockaddr rt_dst; /* key */ struct sockaddr rt_gateway; /* value */ short rt_flags; /* up/down?, host/net */ short rt_refcnt; /* # held references */ unit_t rt_use; /* raw # packets forwarded */ /* * The kernel does not use this field, and without it the structure is * datamodel independent. */ #ifdef !_KERNEL struct ifnet *rt_ifp; /* the answer: interface to use */ #endif };</pre> <p>with <code>rt_flags</code> defined from:</p> <pre>#define RTF_UP 0x1 /* route usable */ #define RTF_GATEWAY 0x2 /* destination is a gateway */ #define RTF_HOST 0x4 /* host entry (net otherwise) */</pre>

is a “direct” connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface may be requested to address the packet to an entity different from the eventual recipient; essentially, the packet is forwarded.

Routing table entries installed by a user process may not specify the hash, reference count, use, or interface fields; these are filled in by the routing routines. If a route is in use when it is deleted, meaning its `rt_refcnt` is non-zero, the resources associated with it will not be reclaimed until all references to it are removed.

User processes read the routing tables through the `/dev/ip` device.

The `rt_use` field contains the number of packets sent along the route. This value is used to select among multiple routes to the same destination. When multiple routes to the same destination exist, the least used route is selected.

A wildcard routing entry is specified with a zero destination address value. Wildcard routes are used only when the system fails to find a route to the destination host and network. The combination of wildcard routes and routing redirects can provide an economical mechanism for routing traffic.

ERRORS	EEXIST	A request was made to duplicate an existing entry.
	ESRCH	A request was made to delete a non-existent entry.
	ENOBUFS	Insufficient resources were available to install a new route.
	ENOMEM	Insufficient resources were available to install a new route.
	ENETUNREACH	The gateway is not directly reachable. For example, it does not match the destination/subnet on any of the network interfaces.
FILES	<code>/dev/ip</code>	IP device driver
SEE ALSO	<code>route(1M)</code> , <code>ioctl(2)</code> , <code>route(7P)</code>	

rtls(7D)

NAME	rtls – Realtek Fast Ethernet 8139–based network interface controllers
SYNOPSIS	<code>/dev/rtls</code>
DESCRIPTION	The <code>rtls</code> Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code> , over Realtek 8139 controllers. The <code>rtls</code> driver provides basic support for the Realtek 8139 hardware. Functions include chip initialization, frame transmit and receive, multicast support, and error recovery and reporting.
APPLICATION PROGRAMMING INTERFACE	The cloning, character-special device <code>/dev/rtls</code> is used to access all 8139 devices installed within the system.
rtls and DLPI	<p>The <code>rtls</code> driver is dependent on <code>/kernel/misc/gld</code>, a loadable kernel module that provides the <code>rtls</code> driver with the DLPI and STREAMS functionality required of a LAN driver. See gld(7D) for more details on the primitives supported by the driver.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> from the user are as follows:</p> <ul style="list-style-type: none">■ Maximum SDU is 1500 (ETHERMTU).■ Minimum SDU is 0. The driver pads to the mandatory 60-octet minimum packet size.■ The <code>dlsap</code> address length is 8.■ MAC type is <code>DL_ETHER</code>.■ The <code>sap</code> length value is -2, meaning the physical address component is followed immediately by a 2-byte <code>sap</code> component within the DLSAP address.■ Broadcast address value is Ethernet/IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>).
KNOWN PROBLEMS AND LIMITATIONS	<p>The <code>rtls</code> driver only supports Realtek 8139 on x86 platforms.</p> <p>The <code>rtls</code> driver does not support the NDD interface and MII statistics. Realtek 8139 works under 100FDX/100HDX and 10FDX/10HDX modes, which you can set using the <code>rtls.conf</code> file.</p>
CONFIGURATION	<p>To configure the <code>rtls</code> driver:</p> <ul style="list-style-type: none">■ Use <code>prtconf -v grep pci10ec,8139</code> to obtain the instance number of the driver. (<code>pci1186,1301</code> or <code>pci1113,1211</code> indicates a third-party ethernet device with an 8139 chip that can also be controlled by the <code>rtls</code> driver.) <p>The <code>rtls.conf</code> configuration file options include:</p> <ul style="list-style-type: none">-ForceSpeedDuplex Default: 5 (Auto-negotiate) <p>Allowed Values: 4 (100 FDX)</p>

3 (100 HDX)

2 (10 FDX)

1 (10 HDX)

Specify the speed and duplex mode for each instance.

Example: ForceSpeedDuplex=5,4;

Sets `rtls0` to autonegotiate and `rtls1` to 100 FDX.

FILES

`rtls`

Device special file.

`/kernel/drv/rtls.conf`

`rtls` configuration file.

`<sys/stropts.h>`

`stropts` network header file.

`<sys/ethernet.h>`

Ethernet network header file.

`<sys/dlpi.h>`

`dlpi` network header file.

`<sys/gld.h>`

`gld` network header file.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO

`attributes(5)`, `dlpi(7P)`, `gld(7D)`

sad(7D)

NAME	sad – STREAMS Administrative Driver								
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/conf.h> #include <sys/sad.h> #include <sys/stropts.h> int ioctl(int <i>fildev</i>, int <i>command</i>, int <i>arg</i>);</pre>								
DESCRIPTION	<p>The STREAMS Administrative Driver provides an interface for applications to perform administrative operations on STREAMS modules and drivers. The interface is provided through <code>ioctl(2)</code> commands. Privileged operations may access the <code>sad</code> driver using <code>/dev/sad/admin</code>. Unprivileged operations may access the <code>sad</code> driver using <code>/dev/sad/user</code>.</p> <p>The <i>fildev</i> argument is an open file descriptor that refers to the <code>sad</code> driver. The <i>command</i> argument determines the control function to be performed as described below. The <i>arg</i> argument represents additional information that is needed by this command. The type of <i>arg</i> depends upon the command, but it is generally an integer or a pointer to a command-specific data structure.</p>								
COMMAND FUNCTIONS	<p>The autopush facility (see <code>autopush(1M)</code>) allows one to configure a list of modules to be automatically pushed on a stream when a driver is first opened. Autopush is controlled by the following commands:</p> <p>SAD_SAP Allows the administrator to configure the given device's autopush information. <i>arg</i> points to a <code>strpush</code> structure, which contains the following members:</p> <pre>unit_t ap_cmd; major_t sap_major; minor_t sap_minor; minor_t sap_lastminor; unit_t sap_npsh; unit_t sap_list [MAXAPUSH] [FMNAMESZ + 1];</pre> <p>The <code>sap_cmd</code> field indicates the type of configuration being done. It may take on one of the following values:</p> <table><tr><td><code>SAP_ONE</code></td><td>Configure one minor device of a driver.</td></tr><tr><td><code>SAP_RANGE</code></td><td>Configure a range of minor devices of a driver.</td></tr><tr><td><code>SAP_ALL</code></td><td>Configure all minor devices of a driver.</td></tr><tr><td><code>SAP_CLEAR</code></td><td>Undo configuration information for a driver.</td></tr></table> <p>The <code>sap_major</code> field is the major device number of the device to be configured. The <code>sap_minor</code> field is the minor device number of the device to be configured. The <code>sap_lastminor</code> field is used only with the <code>SAP_RANGE</code> command, which configures a range of</p>	<code>SAP_ONE</code>	Configure one minor device of a driver.	<code>SAP_RANGE</code>	Configure a range of minor devices of a driver.	<code>SAP_ALL</code>	Configure all minor devices of a driver.	<code>SAP_CLEAR</code>	Undo configuration information for a driver.
<code>SAP_ONE</code>	Configure one minor device of a driver.								
<code>SAP_RANGE</code>	Configure a range of minor devices of a driver.								
<code>SAP_ALL</code>	Configure all minor devices of a driver.								
<code>SAP_CLEAR</code>	Undo configuration information for a driver.								

minor devices between `sap_minor` and `sap_lastminor`, inclusive. The minor fields have no meaning for the `SAP_ALL` command. The `sap_npush` field indicates the number of modules to be automatically pushed when the device is opened. It must be less than or equal to `MAXAPUSH`, defined in `sad.h`. It must also be less than or equal to `NSTRPUSH`, the maximum number of modules that can be pushed on a stream, defined in the kernel master file. The field `sap_list` is an array of NULL-terminated module names to be pushed in the order in which they appear in the list.

When using the `SAP_CLEAR` command, the user sets only `sap_major` and `sap_minor`. This will undo the configuration information for any of the other commands. If a previous entry was configured as `SAP_ALL`, `sap_minor` should be set to zero. If a previous entry was configured as `SAP_RANGE`, `sap_minor` should be set to the lowest minor device number in the range configured.

On failure, `errno` is set to the following value:

<code>EFAULT</code>	<code>arg</code> points outside the allocated address space.
<code>EINVAL</code>	The major device number is invalid, the number of modules is invalid, or the list of module names is invalid.
<code>ENOSTR</code>	The major device number does not represent a STREAMS driver.
<code>EEXIST</code>	The major-minor device pair is already configured.
<code>ERANGE</code>	The command is <code>SAP_RANGE</code> and <code>sap_lastminor</code> is not greater than <code>sap_minor</code> , or the command is <code>SAP_CLEAR</code> and <code>sap_minor</code> is not equal to the first minor in the range.
<code>ENODEV</code>	The command is <code>SAP_CLEAR</code> and the device is not configured for autopush.
<code>ENOSR</code>	An internal autopush data structure cannot be allocated.

`SAD_GAP`

Allows any user to query the `sad` driver to get the autopush configuration information for a given device. `arg` points to a `strapush` structure as described in the previous command.

sad(7D)

The user should set the `sap_major` and `sap_minor` fields of the `strpush` structure to the major and minor device numbers, respectively, of the device in question. On return, the `strpush` structure will be filled in with the entire information used to configure the device. Unused entries in the module list will be zero-filled.

On failure, `errno` is set to one of the following values:

EFAULT	<code>arg</code> points outside the allocated address space.
EINVAL	The major device number is invalid.
ENOSTR	The major device number does not represent a STREAMS driver.
ENODEV	The device is not configured for autopush.

SAD_VML

Allows any user to validate a list of modules (that is, to see if they are installed on the system). `arg` is a pointer to a `str_list` structure with the following members:

```
int    sl_nmods;
struct str_mlist *sl_modlist;

char  l_name[FMNAMESZ+1];
```

The `str_mlist` structure has the following member:

`sl_nmods` indicates the number of entries the user has allocated in the array and `sl_modlist` points to the array of module names. The return value is 0 if the list is valid, 1 if the list contains an invalid module name, or -1 on failure. On failure, `errno` is set to one of the following values:

EFAULT	<code>arg</code> points outside the allocated address space.
EINVAL	The <code>sl_nmods</code> field of the <code>str_list</code> structure is less than or equal to zero.

SEE ALSO `intro(2)`, `ioctl(2)`, `open(2)`

STREAMS Programming Guide

DIAGNOSTICS Unless otherwise specified, the return value from `ioctl()` is 0 upon success and -1 upon failure with `errno` set as indicated.

NAME	sbpro – Sound Blaster Pro, Sound Blaster 16, and Sound Blaster AWE32 audio device driver
SYNOPSIS	<pre>sbpro:sound, sbpro sbpro:sound, sbproctl</pre>
DESCRIPTION	<p>The Creative Labs Sound Blaster family of audio cards comprises DMA-capable ISA bus plug-in cards that provide 8 and 16 bit mono and stereo digitized sound recording and playback over a wide range of sampling rates. Each card includes a digital sound processor and mixing capability. Some of the cards also support more advanced audio features such as FM synthesis, advanced signal processing, advanced wave effects, and MIDI capability; however, the <code>sbpro</code> driver does not currently support those advanced features. The features and interfaces supported by the Solaris <code>sbpro</code> driver are described here and in <code>audio(7I)</code>.</p> <p>Some Sound Blaster cards support optional non-audio capabilities such as SCSI interfaces and CD-ROM interfaces. These interfaces are not supported by the <code>sbpro</code> driver.</p> <p>The <code>sbpro</code> driver also supports certain Sound Blaster-compatible audio devices, including some based on the ESS688 audio chip.</p> <p>In addition, the driver supports some devices based on the Analog Devices AD1847 and AD1848, and Crystal Semiconductor CS4231 chips. Any CS4231-based devices supported by this driver are programmed in AD1848 compatibility mode. There is no special support in this driver for the more advanced CS4231 features. This family of devices will be referred to as the "AD184x family."</p>
APPLICATION PROGRAMMING INTERFACE	<p>The Sound Blaster device is treated as an exclusive resource, meaning that only one process may open the device at a time. Since the Sound Blaster hardware does not support simultaneous sound input and output, the <code>sbpro</code> driver does not allow the simultaneous access of the device by two processes, even if one tries to open it read-only and the other write-only.</p> <p>The <code>sbpro</code> driver will return "SUNW, sbpro" or "SUNW, sb16" in the <code>name</code> field of the <code>audio_device</code> structure. The <code>version</code> field will contain the version number of the card's DSP chip, and the <code>config</code> field will be set to "SBPRO" or "SB16". The AWE32 is currently identified as an SB16. In all subjects covered in this man page, the Sound Blaster AWE32 behaves the same as the Sound Blaster 16.</p>
Audio Data Formats	<p>The Sound Blaster Pro handles 8-bit samples. In mono mode, audio data may be sampled at rates from 4,000 to 44,100 samples per second. In stereo mode, samples may be handled at the rates of 11,025 and 22,050 samples per second. The SB-16 can sample 8-bit or 16-bit mono or stereo data in the range of 5,000 to 44,100 Hz. Devices in the AD184x family can handle sample rates up to 48,000 Hz.</p>

sbpro(7D)

	<p>The Sound Blaster Pro hardware handles 8-bit linear samples in excess-128 format. The Sound Blaster 16 handles that format as well as 16-bit linear samples in two's complement format. The <code>sbpro</code> driver will generate and accept data in these formats if <code>AUDIO_ENCODING_LINEAR</code> is selected in the <i>encoding</i> field of the audio information structure. 16 bit precision is not available on the Sound Blaster Pro. The <code>sbpro</code> driver will also accept and generate mu-law format data (as in the Greek letter mu) if the <i>encoding</i> field is set to <code>AUDIO_ENCODING_ULAW</code>. In this case, driver software performs the translation between linear and mu-law formats. mu-law encoding is designed to provide an improved signal-to-noise ratio at low amplitude levels. To achieve best results when using mu-law encoding, the audio record volume should be set so that typical amplitude levels lie within approximately three-fourths of the full dynamic range. Devices in the AD184x family support both mu-law and A-law in hardware, and the driver allows either of those encodings to be selected.</p>
<p>Audio Ports</p>	<p>The Sound Blaster hardware does not support multiple output devices, so the <i>play.port</i> field of the audio information structure only supports <code>AUDIO_HEADPHONE</code>. Output volume is controlled by software. The volume control thumbwheel on the back of the card should be turned all the way up to maximum; otherwise no sound may be audible.</p> <p>The <i>record.port</i> field of the audio information structure allows selection of which audio source is used for recording, and may be set to one of <code>AUDIO_MICROPHONE</code>, <code>AUDIO_LINE_IN</code>, or <code>AUDIO_CD</code>. These select input from the microphone jack, line-level input jack, or internal CD input, respectively. The microphone input is treated as a mono source by the hardware, although the microphone jack is a stereo jack. If your microphone has a mono plug, you should convert it to a stereo plug using an appropriate adapter. Line and CD are stereo sources. When recording in mono mode, both stereo channels are mixed before recording.</p>
<p>PRECONFIGURATION</p>	<p>Note – Many audio devices come with a software utility that allows you to select the IRQ and DMA settings. Often, this utility does not record parameters in nonvolatile memory but in a configuration file used by DOS to set the card's configuration at each reboot. This type of configuration file is not used by the Solaris software and does not affect the operation of the card with the Solaris operating environment.</p> <ul style="list-style-type: none">■ Output volume is controlled by software. If you do not hear sound, turn the volume thumbwheel on the back of the card to the maximum volume setting.■ Line-in and aux jacks typically require line level voltages, such as output from a tape or CD player line-out jack or from a powered (battery-operated) microphone. Mic jacks typically require lower voltages. Consult the manufacturer's documentation for your device requirements.■ The Sound Blaster Pro card cannot share IRQ settings with any other card installed in your system. If the hardware-jumpered IRQ setting conflicts with any other device, change the IRQ on the Sound Blaster card to one listed under Supported Settings. The most common conflicts occur with the LPT1 parallel port or a network card.

Supported Settings

- For Sound Blaster 16 cards that have an on-board SCSI subsystem, the audio subsystem needs its own I/O (port) address and an IRQ, distinct from those of the SCSI subsystem.

Note – Default settings differ slightly between the SB Pro, SB Pro-2 and SB 16, SB AWE32, SB Vibra 16 cards. Only the 0x220 and 0x240 settings are supported for SB Pro and SB Pro-2 while additional addresses are supported for the SB 16, SB AWE32, and SB Vibra 16 cards.

If your card supports Plug and Play, your device resources are configured automatically. Use the following settings for devices that don't support Plug and Play. Defaults settings are shown in italics.

- IRQ Level: 2, 5, 7, 10
- I/O Address: *0x220, 0x240, 0x260, 0x280*
- 8-bit DMA Channel: 0, 1, 3
- 16-bit DMA Channel: 5, 6, 7

Settings for the Compaq Deskpro XL and Turtle Beach Tropez Card are provided below. For each device, the sbpro driver automatically chooses an unused DMA channel and IRQ line.

- Compaq Deskpro XL Business Audio With Built-in AD1847 Chip — I/O Address: *0x530, 0x604, 0xE80, 0xF40*

Note – The sbpro support for the AD1848 and compatibles uses one DMA channel for both play and record; simultaneous play/record is not supported.

- Turtle Beach Tropez Card With CS4231 Chip — I/O Address: *0x530* (The MWSS I/O address on the Turtle Beach Tropez card is 0x530 at power-up. It can only be changed by software after the system is booted, a function that the Solaris operating environment does not perform. Therefore, the Tropez card is only supported at I/O address 0x530.)

Note – The Tropez card comes with a software utility for selecting the IRQ, DMA, and MWSS compatibility I/O address settings used by the card. However, that utility does not record those parameters in nonvolatile memory, but in a configuration file used by DOS to set the card's configuration at each reboot. This type of configuration file is not used by the Solaris software and does not affect the operation of the card with the Solaris operating environment.

Known Problems and Limitations

- Any Crystal Semiconductor CS4231-based devices supported by this driver are programmed in AD1848-compatibility mode. This driver does not include support for advanced CS4231 features; in particular, simultaneous play/record.
- The Sound Blaster card cannot share IRQ settings with any other card installed in your system. The most common conflicts occur with the LPT1 parallel port or a network card.
- Some devices can detect that the IRQ is in use by another device in the system. If this occurs, the driver prints an error message:

sbpro(7D)

sbpro: MWSS_AD184x IRQ 7 is 'in use.'

To correct this, change the IRQ setting of either the audio device or the conflicting device. (Some devices are not able to detect such a conflict. The driver will try to use the card, but that will likely result in the system hanging when the card is first used. Thus, it is important to check that the IRQ that does not conflict with another device.)

- Although the `sbpro` driver supports A-law encoding on AD1848 and compatible devices, `audiotool` does not and produces an error message if you select A-law encoding. Use `audioplay(1)` to play A-law encoded audio files, or use `audioconvert(1)` to convert the A-law sample into a format that `audiotool` will accept, such as 16-bit linear. User-written applications can select A-law format using the `sbpro` driver on AD1848 and compatible devices.

Note – Some Compaq Deskpro XL Business Audio system units with built-in AD184x chip have the headphone jack wired with its left and right channels reversed. As a result, left and right output is reversed. The line-out jack at the back of the unit works as expected. For optimum sound quality, use external microphone and speakers and not the ones built into the keyboard.

- Non-Plug and Play Sound Blaster 16, Sound Blaster Vibra 16, and Sound Blaster AWE32 cards are recognized as Sound Blaster 16 cards.
- The ISA version IBM Token Ring and compatible adapters will not work in a system that contains a Sound Blaster card that is configured at the default I/O port address (0x220). If possible, move the Sound Blaster card to port address 0x240; otherwise, remove the Sound Blaster device from the system.

FILES	<code>/dev/audio</code>	Linked to <code>s/dev/sound/0</code>
	<code>/dev/audioc1</code>	Linked to <code>/dev/sound/0c1</code>
	<code>/dev/sound/0</code>	First audio device in the system
	<code>/dev/sound/0c1</code>	Audio control for first audio device
	<code>/usr/share/audio</code>	Audio files

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `audioconvert(1)`, `ioc1(2)`, `attributes(5)`, `audio(7I)`, `streamio(7I)`

Creative Labs, Inc. *Sound Blaster Pro User Reference Manual*

BUGS The current driver implementation does not support the A-law encoding mode for Sound Blaster and compatible devices.

sbpro(7D)

The conversion of mu-law to 8-bit linear format for Sound Blaster and compatible devices can cause a loss of precision, resulting in poor sound quality in cases where the original recording level was well below normal. If this occurs while using the Sound Blaster 16 card, `audioconvert(1)` can be used to convert the original mu-law data to 16-bit linear format before play. This will preserve all the precision from the original mu-law sample.

scmi2c(7d)

NAME	scmi2c – Smart Transporter chip device driver
DESCRIPTION	The <code>scmi2c</code> Smart Transporter device driver is a kernel-loadable Solaris device driver for the Sun Microsystems Smart Transporter chip that features Sun Smartcard internal reader support using the I2C bus interface.
FILES	<code>dev/scmi2cn</code> SCM Microsystems Smart Transporter chip device node <code>/platform/sun4u/kernel/drv/sparcv9/scmi2c</code> SCM Microsystems Smart Transporter chip kernel module
SEE ALSO	<code>smartcard(1M)</code> , <code>smartcard(5)</code>

NAME	scsa1394 – SCSI to 1394 bridge driver
SYNOPSIS	unit@GUID
DESCRIPTION	<p>The <code>scsa1394</code> driver is a 1394 target and an SCSI HBA driver that supports 1394 mass storage devices compliant with the <i>Serial Bus Protocol 2 (SBP-2)</i> specification. It supports both bus-powered and self-powered 1394 mass storage devices.</p> <p>The <code>scsa1394</code> nexus driver maps SCSI target driver requests to SBP-2 Operation Request Blocks (ORB's).</p> <p>The <code>scsa1394</code> driver creates a child device info node for each logical unit (LUN) on the mass storage device. The standard Solaris SCSI disk driver is attached to those nodes. Refer to sd(7D).</p> <p>This driver supports multiple LUN devices and creates a separate child device info node for each LUN. All child LUN nodes attach to sd(7D).</p> <p>All 1394 mass storage devices are treated as removable media devices. A 1394 mass storage device can be managed by <code>rmformat(1)</code>. With or without Volume Manager, you can mount, eject, hot remove and hot insert a 1394 mass storage device, as the following sections explain.</p>
Using Volume Management	<p>Mass storage devices are managed by Volume Manager. <code>vold(1M)</code> creates a device nickname which can be listed with <code>eject(1)</code>. The device is mounted using <code>volrmmount(1)</code> under <code>/rmdisk/label</code>.</p> <p>See <code>volrmmount(1)</code> to unmount the device and <code>eject(1)</code> to eject the media. If the device is ejected while it is mounted, <code>vold(1M)</code> unmounts the device before ejecting it. It also kills any active applications that are accessing the device.</p> <p>Hot removing a mass storage device with <code>vold(1M)</code> active fails with a console warning. To hot remove or insert a 1394 storage device, first stop <code>vold(1M)</code> by issuing the command <code>/etc/init.d/volmgt stop</code>. After the device has been removed or inserted, restart <code>vold(1M)</code> by issuing the command <code>/etc/init.d/volmgt start</code>.</p> <p>You can also permanently disable <code>vold</code> for removable devices by commenting out the <code>rm SCSI</code> line in <code>vold.conf</code>. See the <i>System Administration Guide, Volume I</i> and <i>Solaris Common Desktop Environment: User's Guide</i> for details on how to manage a removable device with CDE and Removable Media Manager. See <code>dtfile.1X</code> under CDE for information on how to use Removable Media Manager.</p>
USING mount(1M) AND umount(1M)	<p>Use <code>mount(1M)</code> to mount the device and <code>umount(1M)</code> to unmount the device. Use <code>eject(1)</code> to eject the media. Because <code>vold(1M)</code> is disabled, no <code>vold</code> nicknames can be used.</p> <p>Removing the storage device while it is being accessed or mounted fails with a console warning. To hot remove the storage device from the system, unmount the file system, then kill all applications accessing the device. Next, hot remove the device. A storage device can be hot inserted at any time.</p>

scsa1394(7D)

For a comprehensive listing of (non-bootable) 1394 mass-storage devices that are compatible with this driver, see *www.sun.com/io*.

DEVICE SPECIAL FILES

Block special file names are located in `/dev/dsk`. Raw file names are located in `/dev/rdisk`. Input/output requests to the devices must follow the same restrictions as those for SCSI disks. Refer to [sd\(7D\)](#).

IOCTLS

Refer to [cdio\(7I\)](#) and [dkio\(7I\)](#).

ERRORS

Refer to [sd\(7D\)](#).

FILES

The device special files for the 1394 mass storage device are created like those for a SCSI disk. Refer to [sd\(7D\)](#).

`/dev/dsk/cntndnsn`

Block files

`/dev/rdisk/cntndnsn`

Raw files

`/vol/dev/aliases/rmdisk0`

Symbolic link to the character device for the media in removable drive 0. This is a generic removable media device.

`/kernel/drv/scsa1394`

32-bit x86 ELF kernel module

`/kernel/drv/amd64/scsa1394`

64-bit x86 ELF kernel module

`/kernel/drv/sparcv9/scsa1394`

64-bit SPARC ELF kernel module

ATTRIBUTES

See [attributes\(5\)](#) for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86, PCI-based systems
Availability	SUNWscsa1394

SEE ALSO

[cdrw\(1\)](#), [eject\(1\)](#), [rmformat\(1\)](#), [volrmmount\(1\)](#), [cfgadm_scsi\(1M\)](#), [fdisk\(1M\)](#), [mount\(1M\)](#), [umount\(1M\)](#), [vold\(1M\)](#), [dtfile.1X](#), [scsi\(4\)](#), [attributes\(5\)](#), [hci1394\(7D\)](#), [sd\(7D\)](#), [pcfs\(7FS\)](#), [cdio\(7I\)](#), [dkio\(7I\)](#)

IEEE Std 1394-1995 Standard for a High Performance Serial Bus

ANSI NCITS 325-1998 - Serial Bus Protocol 2 (SBP-2)

System Administration Guide: Devices and File Systems

Solaris Common Desktop Environment: User's Guide

<http://www.sun.com/io>

scsa2usb(7D)

NAME	scsa2usb – SCSI to USB bridge driver
SYNOPSIS	storage@unit-address
DESCRIPTION	<p>The <code>scsa2usb</code> driver is a USB A (Solaris USB architecture) compliant nexus driver that supports the <i>USB Mass Storage Bulk Only Transport Specification 1.0</i> and <i>USB Control/Bulk/Interrupt (CBI) Transport Specification 1.0</i>. The <code>scsa2usb</code> driver also supports USB storage devices that implement CBI Transport without the interrupt completion for status (that is, Control/Bulk (CB) devices.) It supports bus-powered and self-powered USB mass storage devices. This nexus driver is both a USB client driver and a SCSI HBA driver. As such, the <code>scsa2usb</code> driver only supports storage devices that utilize the above two transports.</p> <p>The <code>scsa2usb</code> nexus driver maps SCSI target driver requests to USB A client driver requests.</p> <p>The <code>scsa2usb</code> driver creates a child device info node for each logical unit (LUN) on the mass storage device. The standard Solaris SCSI disk driver or tape driver is attached to those nodes. Refer to sd(7D) or st(7D).</p> <p>This driver supports multiple LUN devices and creates a separate child device info node for each LUN. All child LUN nodes attach to sd(7D) for disks or st(7D) for tapes.</p> <p>All USB disk storage devices are treated as removable media devices. A USB disk storage device can be managed by <code>rmformat(1)</code>. With or without Volume Manager, you can mount, eject, hot remove and hot insert a USB disk storage device, as the following sections explain.</p> <p>Some devices may be supported by the USB mass storage driver even though they do not identify themselves as compliant with the USB mass storage class.</p> <p>The <code>scsa2usb.conf</code> file contains an <code>attribute-override-list</code> that lists the vendor ID, product ID, and revision for matching mass storage devices, as well as fields for overriding the default device attributes. The entries in this list are commented out by default and may be uncommented to enable support of particular devices.</p> <p>Follow the information given in the <code>scsa2usb.conf</code> file to see if a particular device can be supported using the override information. Also see http://www.sun.com/io. For example, by adding the following to the <code>scsa2usb.conf</code> file, many USB memory sticks and card readers might operate more reliably:</p> <pre>attribute-override-list = "vid=* reduced-cmd-support=true";</pre> <p>Note that this override applies to all USB mass storage devices and might be inappropriate for a USB CD writer. If so, you can add an entry for each device to the attribute override list.</p> <p>If USB mass storage support is considered a security risk, this driver can be disabled in <code>/etc/system</code> as follows:</p>

Using Volume Management	<pre>exclude: scsa2usb</pre> <p>Alternatively, you can permanently disable <code>vold(1M)</code>.</p> <p>Disk storage devices are managed by Volume Manager. <code>vold(1M)</code> creates a device nickname which can be listed with <code>eject(1)</code>. The device is mounted using <code>volrmmount(1)</code> under <code>/rmdisk/label</code>.</p> <p>See <code>volrmmount(1)</code> to unmount the device and <code>eject(1)</code> to eject the media. If the device is ejected while it is mounted, <code>vold(1M)</code> unmounts the device before ejecting it. It also kills any active applications that are accessing the device.</p> <p>Hot removing a disk storage device with <code>vold(1M)</code> active fails with a console warning. To hot remove or insert a USB disk storage device, first stop <code>vold(1M)</code> by issuing the command <code>/etc/init.d/volmgt stop</code>. After the device has been removed or inserted, restart <code>vold(1M)</code> by issuing the command <code>/etc/init.d/volmgt start</code>.</p> <p>You can also permanently disable <code>vold</code> for removable devices by commenting out the <code>rm SCSI</code> line in <code>vold.conf</code>. See the <i>System Administration Guide, Volume I</i> and <i>Solaris Common Desktop Environment: User's Guide</i> for details on how to manage a removable device with CDE and Removable Media Manager. See <code>dtfile.1X</code> under CDE for information on how to use Removable Media Manager.</p>
Using mount(1M) and umount(1M)	<p>Use <code>mount(1M)</code> to mount disk devices and <code>umount(1M)</code> to unmount disk devices. Use <code>eject(1)</code> to eject the media. No <code>vold</code> nicknames can be used. (<code>vold.1m</code> is disabled.)</p> <p>Removing the disk device while it is being accessed or mounted fails with a console warning. To hot remove the disk device from the system, unmount the file system, then kill all applications accessing the device. Next, hot remove the device. A storage device can be hot inserted at any time.</p> <p>For a comprehensive listing of (non-bootable) USB mass-storage devices that are compatible with this driver, see www.sun.com/io.</p>
DEVICE SPECIAL FILES	<p>Disk block special file names are located in <code>/dev/dsk</code>, while raw file names are located in <code>/dev/rdsk</code>. Tape raw file names are located in <code>/dev/rmt</code>. Input/output requests to the devices must follow the same restrictions as those for SCSI disks or tapes. Refer to sd(7D) or st(7D).</p>
IOCTLS	<p>Refer to dkio(7I) and cdio(7I).</p>
ERRORS	<p>Refer to sd(7D) for disks or st(7D) for tapes.</p>
FILES	<p>The device special files for the USB mass storage device are created like those for a SCSI disk or SCSI tape. Refer to sd(7D) or st(7D).</p> <pre>/dev/dsk/cntndnsn Block files for disks.</pre> <pre>/dev/rdsk/cntndnsn Raw files for disks.</pre>

scsa2usb(7D)

/dev/rmt/[0-127] [l,m,h,u,c] [b] [n]
Raw files for tapes.

/vol/dev/aliases/zip0
Symbolic link to the character device for the media in Zip drive 0.

/vol/dev/aliases/jaz0
Symbolic link to the character device for the media in Jaz drive 0.

/vol/dev/aliases/rmdisk0
Symbolic link to the character device for the media in removable drive 0. This is a generic removable media device.

/kernel/drv/scsa2usb
32-bit x86 ELF kernel module

/kernel/drv/amd64/scsa2usb
64-bit x86 ELF kernel module

/kernel/drv/sparcv9/scsa2usb
64-bit SPARC ELF kernel module

/kernel/drv/scsa2usb.conf
Can be used to override specific characteristics.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC, x86, PCI-based systems
Availability	SUNWusb

SEE ALSO `cdrw(1)`, `eject(1)`, `rmformat(1)`, `volrmmount(1)`, `cfgadm_scsi(1M)`, `cfgadm_usb(1M)`, `fdisk(1M)`, `mount(1M)`, `umount(1M)`, `vold(1M)`, `dtfile.1X`, `scsi(4)`, `attributes(5)`, `sd(7D)`, `st(7D)`, `usba(7D)`, `pcfs(7FS)`, `cdio(7I)`, `dkio(7I)`

Writing Device Drivers

System Administration Guide, Volume I

Solaris Common Desktop Environment: User's Guide

Universal Serial Bus Specification 2.0

Universal Serial Bus Mass Storage Class Specification Overview 1.0

Universal Serial Bus Mass Storage Class Bulk-Only Transport Specification 1.0

Universal Serial Bus Mass Storage Class Control/Bulk/Interrupt (CBI) Transport Specification 1.0

System Administration Guide: Basic Administration

<http://www.sun.com/io>

DIAGNOSTICS

Refer to [sd\(7D\)](#) and [st\(7D\)](#).

In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

```
Warning: <device path> (scsa2usb<instance number>): Error Message...
```

Cannot access <device>. Please reconnect.

There was an error in accessing the mass-storage device during reconnect. Please reconnect the device.

Device is not identical to the previous one on this port. Please disconnect and reconnect.

Another USB device has been inserted on a port that was connected to a mass-storage device. Please disconnect the USB device and reconnect the mass-storage device back into that port.

Disconnected device was busy, please reconnect.

Disconnection of the mass-storage device failed because the device is busy. Please reconnect the device.

Reinserted device is accessible again.

The mass-storage device that was hot-removed from its USB slot has been re-inserted to the same slot and is available for access.

Please disconnect and reconnect this device.

A hotplug of the device is needed before it can be restored.

NOTES

The Zip 100 drive does not comply with *Universal Serial Bus Specification 1.0* and cannot be power managed. Power Management support for Zip 100 has been disabled.

If the system panics while a UFS file system is mounted on the mass storage media, no syncing will take place for the disk mass-storage device. (Syncing is not supported by the `scsa2usb` driver.) As a result, the file system on the media will not be consistent on reboot.

If a PCFS file system is mounted, no syncing is needed and the filesystem will be consistent on reboot.

If a mass-storage device is busy, system suspend cannot proceed and the system will immediately resume again.

Attempts to remove a mass-storage device from the system will fail. The failure will be logged to the console. An attempt to replace the removed device with some other USB device will also fail. To successfully remove a USB mass-storage device you must "close" all references to it.

An Iomega Zip 100Mb disk cannot be formatted on an Iomega Zip250 drive. See the Iomega web site at <http://www.iomega.com> for details.

scsa2usb(7D)

Concurrent I/O to devices with multiple LUNs on the same device is not supported.

Some USB CD-RW devices may perform inadequately at their advertised speeds. To compensate, use USB CD-RW devices at lower speeds (2X versus 4X). See `cdrw(1)` for details.

This driver also supports CBI devices that do not use USB interrupt pipe for status completion.

NAME	scsi_vhci – SCSI virtual host controller interconnect driver
DESCRIPTION	<p>The <code>scsi_vhci</code> driver is a SCSI compliant pseudo nexus driver that supports Solaris operating system I/O multipathing services for SCSI-3 devices. This driver introduces a fundamental restructuring of the Solaris device tree to enable a multipath device to be represented as single device instance rather than as an instance per physical path as in earlier Solaris versions.</p> <p>The multipath SCSI target devices managed by this driver are identified and represented by using SCSI-3 VPD page (0x83) as the device's unit address.</p> <p>Symbolic links in <code>/dev/[r]dsk</code> continue to adhere to the <code>cNtNdNsN</code> format. <code>cN</code> is the logical controller number assigned to this driver instance. <code>tN</code> is the global unique identifier (GUID) of the multipath target device (64/128 bits), represented as hexadecimal numbers.</p> <p>The following is an example of a system with a A5000 storage array:</p> <pre>... /dev/rdisk/c4t200000203709C3F5d0s0 -> ../../devices/scsi_vhci/ ssd@g200000203709c3f5:a,raw ... /dev/rdisk/c4t200000203709C3F5d0s7 -> ../../devices/scsi_vhci/ ssd@g200000203709c3f5:h,ra ...</pre> <p>The following is an example of a system with a T300 storage array:</p> <pre>... /dev/rdisk/c1t60020F200000033939C2C2B60008D4AEd0s0 -> ../../devices/ scsi_vhci/ssd@g60020f200000033939a2c2b60008d4ae:a,raw ... /dev/rdisk/c1t60020F200000033939A2C2B60008D4AEd0s7 -> ../../devices/ scsi_vhci/ssd@g60020f200000033939a2c2b60008d4ae:h,raw ...</pre> <p>The <code>scsi_vhci</code> driver receives naming and transport services from one or more physical HBA (host bus adapter) devices. To support multi-pathing, a physical HBA driver must comply with the multipathing services provided by this driver.</p> <p>The <code>scsi_vhci</code> driver supports the standard functions provided by the SCSI interface.</p>
Configuration	<p>The <code>scsi_vhci</code> driver can be configured by defining properties in the <code>scsi_vhci.conf</code> file.</p> <p>The <code>scsi_vhci</code> driver supports the following property:</p> <pre>device-type-scsi-options-list</pre> <p>To add a third-party (non-Sun) symmetric storage device to run under <code>scsi_vhci</code> (and thereby take advantage of Solaris I/O multipathing), you add the vendor ID and product ID for the device, as those strings are returned by the SCSI Inquiry command. As shipped, the <code>scsi_vhci.conf</code> file contains, in part:</p>

scsi_vhci(7D)

```
# device-type-scsi-options-list =  
# "SUN      SENA", "symmetric-option";  
# symmetric-option=0x1000000;
```

For example, to add a device from a vendor with the ID of "Acme" and a product ID of "MSU", you would add:

```
device-type-scsi-options-list =  
"Acme    MSU", "symmetric-option";  
symmetric-option=0x1000000;
```

In addition to "Acme," you also might want to add another entry, for example, a device from "XYZ" vendor with a product ID of "ABC:"

```
device-type-scsi-options-list =  
"Acme    MSU", "symmetric-option",  
"XYZ     ABC", "symmetric-option";  
symmetric-option=0x1000000;
```

FILES /kernel/drv/sparcv9/scsi_vhci
64-bit kernel module (SPARC).

/kernel/drv/scsi_vhci
32-bit kernel module (x86).

/kernel/drv/scsi_vhci.conf
Driver configuration file.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	PCI-based systems
Availability	SUNWckr

SEE ALSO [eeprom\(1M\)](#), [prtconf\(1M\)](#), [attributes\(5\)](#), [fcp\(7D\)](#), [fctl\(7D\)](#), [fp\(7D\)](#), [ssd\(7D\)](#), [scsi_abort\(9F\)](#), [scsi_ifgetcap\(9F\)](#), [scsi_pkt\(9S\)](#), [scsi_reset\(9F\)](#), [scsi_transport\(9F\)](#), [scsi_inquiry\(9S\)](#), [scsi_extended_sense\(9S\)](#)

Writing Device Drivers

Small Computer System Interface-3 (SCSI-3)

NOTES In previous releases, the `scsi_vhci.conf` file supported the `mpxio-disable` property, which allowed you to disable Solaris I/O multipathing on a system-wide basis. This property is not present in the current release of the Solaris operating system. Multipathing is always enabled in `scsi_vhci`. If you want to disable multipathing, use the mechanisms provided by the HBA drivers. See [fp\(7D\)](#).

In previous releases, Solaris I/O multipathing was also known as MPxIO and Sun StorEdge Traffic Manager (STMS).

NAME	sctp, SCTP – Stream Control Transmission Protocol
SYNOPSIS	<pre>#include <sys/socket.h> #include <netinet/in.h> s = socket(AF_INET, SOCK_STREAM, IPPROTO_SCTP); s = socket(AF_INET, SOCK_SEQPACKET, IPPROTO_SCTP); s = socket(AF_INET6, SOCK_STREAM, IPPROTO_SCTP); s = socket(AF_INET6, SOCK_SEQPACKET, IPPROTO_SCTP);</pre>
DESCRIPTION	<p>SCTP is a transport protocol layered above the Internet Protocol (IP), or the Internet Protocol Version 6 (IPv6). SCTP provides a reliable, session oriented, flow-controlled, two-way transmission of data. It is a message-oriented protocol and supports framing of individual messages boundaries. An SCTP association is created between two endpoints for data transfer which is maintained during the lifetime of the transfer. An SCTP association is setup between two endpoints using a four-way handshake mechanism with the use of a cookie to guard against some types of denial of service (DoS) attacks. These endpoints may be represented by multiple IP addresses.</p> <p>An SCTP message includes a common SCTP header followed by one or more chunks. Included in the common header is a 32-bit field which contains the checksum (computed using CRC-32c polynomial) of the entire SCTP packet.</p> <p>SCTP transfers data payloads in the form of DATA chunks. Each DATA chunk contains a Transmission Sequence Number (TSN), which governs the transmission of messages and detection of loss. DATA chunk exchanges follow the Transmission Control Protocol's (TCP) Selective ACK (SACK) mechanism. The receiver acknowledges data by sending SACK chunks, which not only indicate the cumulative TSN range received, but also non-cumulative TSNs received, implying gaps in the received TSN sequence. SACKs are sent using the delayed ack method similar to TCP, that is, one SCTP per every other received packet with an upper bound on the delay (when there are gaps detected the frequency is increased to one every received packet). Flow and congestion control follow TCP algorithms: Slow Start, Congestion Avoidance, Fast Recovery and Fast retransmit. But unlike TCP, SCTP does not support half-close connection and "urgent" data.</p> <p>SCTP is designed to support a number of functions that are critical for telephony signalling transport, including multi-streaming. SCTP allows data to be partitioned into multiple streams that have the property of independent sequenced delivery so that message loss in any one stream only affects delivery within that stream. In many applications (particularly telephony signalling), it is only necessary to maintain sequencing of messages that affect some resource. Other messages may be delivered without having to maintain overall sequence integrity. A DATA chunk on an SCTP association contains the Stream Id/Stream Sequence Number pair, in addition to the TSN, which is used for sequenced delivery within a stream.</p> <p>SCTP uses IP's host level addressing and adds its own per-host collection of port addresses. The endpoints of an SCTP association are identified by the combination of IP address(es) and an SCTP port number. By providing the ability for an endpoint to</p>

sctp(7P)

have multiple IP addresses, SCTP supports multi-homing, which makes an SCTP association more resilient in the presence of network failures (assuming the network is architected to provide redundancy). For a multi-homed SCTP association, a single address is used as the primary address, which is used as the destination address for normal DATA chunk transfers. Retransmitted DATA chunks are sent over alternate address(es) to increase the probability of reaching the remote endpoint. Continued failure to send DATA chunks over the primary address results in selecting an alternate address as the primary address. Additionally, SCTP monitors the reachability of all alternate addresses by sending periodic "heartbeats" chunks. An SCTP association supports multi-homing by exchanging the available list of addresses during association setup (as part of its four-way handshake mechanism). An SCTP endpoint is associated with a local address using the `bind(3SOCKET)` call. Subsequently, the endpoint can be associated with additional addresses using `sctp_bindx(3SOCKET)`. By using a special value of `INADDR_ANY` with IP or the unspecified address (all zeros) with IPv6 in the `bind(3SOCKET)` or `sctp_bindx(3SOCKET)` calls, an endpoint can be bound to all available IP or IPv6 addresses on the system.

SCTP uses a 3-way mechanism to allow graceful shutdown, where each endpoint has confirmation of the DATA chunks received by the remote endpoint prior to completion of the shutdown. An Abort is provided for error cases when an immediate shutdown is needed.

Applications can access SCTP using the socket interface as a `SOCK_STREAM` (one-to-one style) or `SOCK_SEQPACKET` (one-to-many style) socket type.

One-to-one style socket interface supports similar semantics as sockets for connection oriented protocols, such as TCP. Thus, a passive socket is created by calling the `listen(3SOCKET)` system call after binding the socket using `bind(3SOCKET)`. Associations to this passive socket can be received using `accept(3SOCKET)` system call. Active sockets use the `connect(3SOCKET)` system call after binding to initiate an association. If an active socket is not explicitly bound, an implicit binding is performed. If an application wants to exchange data during the association setup phase, it should not call `connect(3SOCKET)`, but use `sendto(3SOCKET)/sendmsg(3SOCKET)` to implicitly initiate an association. Once an association has been established, `read(2)` and `write(2)` system calls can be used to exchange data. Additionally, `send(3SOCKET)`, `recv(3SOCKET)`, `sendto(3SOCKET)`, `recvfrom(3SOCKET)`, `sendmsg(3SOCKET)` and `recvmsg(3SOCKET)` can be used.

One-to-many socket interface supports similar semantics as sockets for connection less protocols, such as UDP (however, unlike UDP, it does not support broadcast or multicast communications). A passive socket is created using the `listen(3SOCKET)` system call after binding the socket using `bind(3SOCKET)`. An `accept(3SOCKET)` call is not needed to receive associations to this passive socket (in fact an `accept(3SOCKET)` on a one-to-many socket will fail). Associations are accepted automatically and notifications of new associations are delivered in `recvmsg(3SOCKET)` provided notifications are enabled. Active sockets after binding (implicitly or explicitly) need not call `connect(3SOCKET)` to establish an association, implicit associations can be created using `sendmsg(3SOCKET)/recvmsg(3SOCKET)`

or `sendto(3SOCKET)/recvfrom(3SOCKET)` calls. Such implicit associations cannot be created using `send(3SOCKET)` and `recv(3SOCKET)` calls. A one-to-one style association can be created from a one-to-many association by branching it off using the `sctp_peeloff(3SOCKET)` call; `send(3SOCKET)` and `recv(3SOCKET)` can be used on such peeled off associations. A `close(2)` on a one-to-many socket will gracefully shutdown all the associations represented by that one-to-many socket.

`sctp_sendmsg(3SOCKET)` and `sctp_recvmsg(3SOCKET)` can be used to avail of advanced features provided by SCTP.

SCTP provides socket option which are set using `setsockopt(3SOCKET)` and read using `getsockopt(3SOCKET)`. The option level is the protocol number for SCTP, available from `getprotobyname(3SOCKET)`. Socket option `SCTP_NODELAY` is used to turn on/off any Nagle like algorithm (similar to `TCP_NODELAY`). Options `SO_RCVBUF` and `SO_SNDBUF` are used to set the receive and send buffer respectively. For one-to-many style socket, `SCTP_AUTOCLOSE` option can be used to automatically close any association that has been idle for more than the specified number of seconds. A value of '0' indicates that no associations should be closed automatically. Option `SCTP_EVENTS` is used to specify various notifications and ancillary data the user wishes to receive. SCTP also provides an option - `SCTP_STATUS` - to retrieve current status information about an SCTP association.

Multihoming

The ability of SCTP to use multiple addresses in an association can create issues with some network utilities. This requires a system administrator to be careful in setting up the system.

For example, the `tcpd(1M)` allows an administrator to use a simple form of address/hostname access control. While `tcpd(1M)` can work with SCTP, the access control part can have some problems. The `tcpd(1M)` access control is only based on one of the addresses at association setup time. Once an association is allowed, no more checking is performed. This means that during the life time of the association, SCTP packets from different addresses of the peer host can be received in the system. This may not be what the system administrator wants as some of the peer's addresses are supposed to be blocked.

Another example is the use of IP Filter, `ipfilter(1M)`. IP Filter provides several functionalities, such as IP packet filtering (`ipf(1M)`) and NAT (`ipnat(1M)`). For packet filtering, one issue is that a filter policy can block packets from some of the addresses of an association while allowing packets from other addresses to go through. This can degrade SCTP's performance when failure occurs. There is a more serious issue with IP address rewrite by NAT. At association setup time, SCTP endpoints exchange IP addresses. But IP Filter is not aware of this. So when NAT is done on a packet, it may change the address to an unacceptable one. Thus the SCTP association setup may succeed but packets cannot go through afterwards when a different IP address is used for the association.

SEE ALSO

`ipfilter(1M)`, `ipf(1M)`, `ipnat(1M)`, `ndd(1M)`, `tcpd(1M)`, `ioctl(2)`, `read(2)`, `write(2)`, `libsctp(3LIB)`, `accept(3SOCKET)`, `bind(3SOCKET)`, `connect(3SOCKET)`, `getprotobyname(3SOCKET)`, `getsockopt(3SOCKET)`,

sctp(7P)

`listen(3SOCKET)`, `recv(3SOCKET)`, `recvfrom(3SOCKET)`, `recvmsg(3SOCKET)`,
`sctp_bindx(3SOCKET)`, `sctp_getladdrs(3SOCKET)`,
`sctp_getpaddrs(3SOCKET)`, `sctp_freeaddrs(1M)`,
`sctp_freepaddrs(3SOCKET)`, `sctp_opt_info(3SOCKET)`,
`sctp_peeloff(3SOCKET)`, `sctp_recvmsg(3SOCKET)`, `sctp_sendmsg(3SOCKET)`,
`send(3SOCKET)`, `sendmsg(3SOCKET)`, `sendto(3SOCKET)`, `socket(3SOCKET)`,
[tcp\(7P\)](#), [udp\(7P\)](#), [inet\(7P\)](#), [inet6\(7P\)](#), [ip\(7P\)](#), [ip6\(7P\)](#)

R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zang, V. Paxson, *RFC 2960, Stream Control Transmission Protocol*, October 2000

L. Ong, J. Yoakum, *RFC 3286, An Introduction to Stream Control Transmission Protocol (SCTP)*, May 2002

J. Stone, R. Stewart, D. Otis, *RFC 3309, Stream Control Transmission Protocol (SCTP) Checksum Change*, September 2002.

DIAGNOSTICS

A socket operation may fail if:

<code>EPROTONOSUPPORT</code>	If the socket type is other than <code>SOCK_STREAM</code> and <code>SOCK_SEQPACKET</code>
<code>ETIMEDOUT</code>	An association was dropped due to excessive retransmissions.
<code>ECONNREFUSED</code>	The remote peer refused establishing an association.
<code>ECONNREFUSED</code>	
<code>EADDRINUSE</code>	A <code>bind()</code> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
<code>EINVAL</code>	A <code>bind()</code> operation was attempted on a socket with an invalid network address.
<code>EPERM</code>	A <code>bind()</code> operation was attempted on a socket with a "reserved" port number and the effective user ID of the process was not the privileged user.

NAME	sd – SCSI disk and ATAPI/SCSI CD-ROM device driver
SYNOPSIS	sd@target , lun:partition
DESCRIPTION	
SPARC	<p>The <code>sd</code> SCSI and SCSI/ATAPI driver supports embedded SCSI-2 and CCS-compatible SCSI disk and CD-ROM drives, ATAPI 2.6 (SFF-8020i)-compliant CD-ROM drives, SFF-8090-compliant SCSI/ATAPI DVD-ROM drives, IOMEGA SCSI/ATAPI ZIP drives, SCSI JAZ drives, and USB mass storage devices (refer to scsa2usb(7D)).</p> <p>To determine the disk drive type, use the SCSI/ATAPI inquiry command and read the volume label stored on block 0 of the drive. (The volume label describes the disk geometry and partitioning and must be present for the disk to be mounted by the system.) A volume label is not required for removable, rewritable or read-only media.</p>
x86 Only	<p>The <code>sd</code> driver supports embedded SCSI-2 and CCS-compatible SCSI disk and CD-ROM drives, ATAPI 2.6 (SFF-8020i)-compliant CD-ROM drives, SFF-8090-compliant SCSI/ATAPI DVD-ROM drives, IOMEGA SCSI/ATAPI ZIP drives, and SCSI JAZ drives.</p> <p>The x86 BIOS legacy requires a master boot record (MBR) and <code>fdisk</code> table in the first physical sector of the bootable media. If the x86 hard disk contains a Solaris disk label, it is located in the second 512-byte sector of the FDISK partition.</p>
DEVICE SPECIAL FILES	<p>Block-files access the disk using normal buffering mechanism and are read-from and written-to without regard to physical disk records. A "raw" interface enables direct transmission between the disk and the user's read or write buffer. A single <code>read</code> or <code>write</code> call usually results in a single I/O operation; raw I/O is therefore more efficient when many bytes are transmitted. Block files names are found in <code>/dev/dsk</code>; raw file names are found in <code>/dev/rdisk</code>.</p> <p>I/O requests to the raw device must be aligned on a 512-byte (<code>DEV_BSIZE</code>) boundary and all I/O request lengths must be in multiples of 512 bytes. Requests that do not meet these requirements will trigger an <code>EINVAL</code> error. There are no alignment or length restrictions on I/O requests to the block device.</p>
CD-ROM DRIVE SUPPORT	<p>A CD-ROM disk is single-sided and contains approximately 640 megabytes of data or 74 minutes of audio. When the CD-ROM is opened, the eject button is disabled to prevent manual removal of the disk until the last <code>close()</code> is called. No volume label is required for a CD-ROM. The disk geometry and partitioning information are constant and never change. If the CD-ROM contains data recorded in a Solaris-aware file system format, it can be mounted using the appropriate Solaris file system support.</p>

sd(7D)

**DVD-ROM
DRIVE SUPPORT**

DVD-ROM media can be single or double-sided and can be recorded upon using a single or double layer structure. Double-layer media provides parallel or opposite track paths. A DVD-ROM can hold from between 4.5 Gbytes and 17 Gbytes of data, depending on the layer structure used for recording and if the DVD-ROM is single or double-sided.

When the DVD-ROM is opened, the eject button is disabled to prevent the manual removal of a disk until the last `close()` is called. No volume label is required for a DVD-ROM. If the DVD-ROM contains data recorded in a Solaris-aware file system format, it can be mounted using the appropriate Solaris file system support.

**ZIP/JAZ DRIVE
SUPPORT**

ZIP/JAZ media provide varied data capacity points; a single JAZ drive can store up to 2 GBytes of data, while a ZIP-250 can store up to 250MBytes of data. ZIP/JAZ drives can be read-from or written-to using the appropriate drive.

When a ZIP/JAZ drive is opened, the eject button is disabled to prevent the manual removal of a disk until the last `close()` is called. No volume label is required for a ZIP/JAZ drive. If the ZIP/JAZ drive contains data recorded in a Solaris-aware file system format, it can be mounted using the appropriate Solaris file system support.

**DEVICE
STATISTICS
SUPPORT**

Each device maintains I/O statistics for the device and for partitions allocated for that device. For each device/partition, the driver accumulates reads, writes, bytes read, and bytes written. The driver also initiates hi-resolution time stamps at queue entry and exit points to enable monitoring of residence time and cumulative residence-length product for each queue.

Not all device drivers make per-partition IO statistics available for reporting. `sd` and `ssd(7D)` per-partition statistics are enabled by default but may be disabled in their configuration files.

**IOCTLS
ERRORS**

Refer to `dkio(7I)`, and `cdio(7I)`

EACCES	Permission denied
EBUSY	The partition was opened exclusively by another thread
EFAULT	The argument features a bad address
EINVAL	Invalid argument
ENOTTY	The device does not support the requested ioctl function
ENXIO	During opening, the device did not exist. During close, the drive unlock failed
EROFS	The device is read-only
EAGAIN	Resource temporarily unavailable
EINTR	A signal was caught during the execution of the <code>ioctl()</code> function

CONFIGURATION

ENOMEM	Insufficient memory
EPERM	Insufficient access permission
EIO	An I/O error occurred. Refer to notes for details on copy-protected DVD-ROM media.

The `sd` driver can be configured by defining properties in the `sd.conf` file. The `sd` driver supports the following properties:

enable-partition-kstats

The default value is 1, which causes partition IO statistics to be maintained. Set this value to zero to prevent the driver from recording partition statistics. This slightly reduces the CPU overhead for IO, mimimizes the amount of `sar(1)` data collected and makes these statistics unavailable for reporting by `iostat(1M)` even though the `-p/-P` option is specified. Regardless of this setting, disk IO statistics are always maintained.

qfull-retries

The supplied value is passed as the `qfull-retries` capability value of the HBA driver. See `scsi_ifsetcap(9F)` for details.

qfull-retry-interval

The supplied value is passed as the `qfull-retry interval` capability value of the HBA driver. See `scsi_ifsetcap(9F)` for details.

allow-bus-device-reset

The default value is 1, which allows resetting to occur. Set this value to 0 (zero) to prevent the `sd` driver from calling `scsi_reset(9F)` with a second argument of `RESET_TARGET` when in error-recovery mode. This `scsi_reset(9F)` call may prompt the HBA driver to send a SCSI Bus Device Reset message. The `scsi_reset(9F)` call with a second argument of `RESET_TARGET` may result from an explicit request via the `USCSICMD ioctl`. Some high-availability multi-initiator systems may wish to prohibit the Bus Device Reset message; to do this, set the `allow-bus-device-reset` property to 0.

FILES

<code>/kernel/drv/sd.conf</code>	driver configuration file
<code>/dev/dsk/cntndnsn</code>	block files
<code>/dev/rdisk/cntndnsn</code>	raw files

Where:

<code>cn</code>	controller <code>n</code>
<code>tn</code>	SCSI target id <code>n</code> (0-6)
<code>dn</code>	SCSI LUN <code>n</code> (0-7 normally; some HBAs support LUNs to 15 or 32. See the specific manpage for details)
<code>sn</code>	partition <code>n</code> (0-7)

sd(7D)

x86 Only /dev/rdisk/cntndnpn raw files

Where:

pn Where $n=0$ the node corresponds to the entire disk.

SEE ALSO sar(1), cfgadm_scsi(1M), fdisk(1M), format(1M), iostat(1M), close(2), ioctl(2), lseek(2), read(2), write(2), driver.conf(4), scsi(4), filesystem(5), scsa2usb(7D), ssd(7D), hfs(7FS), pcfs(7FS), udfs(7FS), cdio(7I), dkio(7I), scsi_ifsetcap(9F), scsi_reset(9F)

ANSI Small Computer System Interface-2 (SCSI-2)

ATA Packet Interface for CD-ROMs, SFF-8020i

Mt.Fuji Commands for CD and DVD, SFF8090v3

<http://www.sun.com/io>

DIAGNOSTICS

Error for Command: '*<command name>*'

Error Level: Fatal

Requested Block: *<n>*

Error Block: *<m>*

Vendor: '*<vendorname>*'

Serial Number: '*<serial number>*'

Sense Key: *<sense key name>*

ASC: 0x*<a>* (*<ASC name>*), ASCQ: 0x**, FRU: 0x*<c>*

The command indicated by *<command name>* failed. The Requested Block is the block where the transfer started and the Error Block is the block that caused the error. Sense Key, ASC, and ASCQ information is returned by the target in response to a request sense command.

Caddy not inserted in drive

The drive is not ready because no caddy has been inserted.

Check Condition on REQUEST SENSE

A REQUEST SENSE command completed with a check condition. The original command will be retried a number of times.

Label says *<m>* blocks Drive says *<n>* blocks

There is a discrepancy between the label and what the drive returned on the READ CAPACITY command.

Not enough sense information

The request sense data was less than expected.

Request Sense couldn't get sense data

The REQUEST SENSE command did not transfer any data.

Reservation Conflict

The drive was reserved by another initiator.

SCSI transport failed: reason 'xxxx': {retrying | giving up}
 The host adapter has failed to transport a command to the target for the reason stated. The driver will either retry the command or, ultimately, give up.

Unhandled Sense Key<n>
 The REQUEST SENSE data included an invalid sense.

Unit not ready. Additional sense code 0x
 <n> The drive is not ready.

Can't do switch back to mode 1
 A failure to switch back to read mode 1.

Corrupt label - bad geometry
 The disk label is corrupted.

Corrupt label - label checksum failed
 The disk label is corrupted.

Corrupt label - wrong magic number
 The disk label is corrupted.

Device busy too long
 The drive returned busy during a number of retries.

Disk not responding to selection
 The drive is powered down or died

Failed to handle UA
 A retry on a Unit Attention condition failed.

I/O to invalid geometry
 The geometry of the drive could not be established.

Incomplete read/write - retrying/giving up
 There was a residue after the command completed normally.

No bp for direct access device format geometry
 A bp with consistent memory could not be allocated.

No bp for disk label
 A bp with consistent memory could not be allocated.

No bp for fdisk
 A bp with consistent memory could not be allocated.

No bp for rigid disk geometry
 A bp with consistent memory could not be allocated.

No mem for property
 Free memory pool exhausted.

No memory for direct access device format geometry
 Free memory pool exhausted.

No memory for disk label
 Free memory pool exhausted.

sd(7D)

No memory for rigid disk geometry
The disk label is corrupted.

No resources for dumping
A packet could not be allocated during dumping.

Offline
Drive went offline; probably powered down.

Requeue of command fails
Driver attempted to retry a command and experienced a transport error.

sdrestart transport failed()
Driver attempted to retry a command and experienced a transport error.

Transfer length not modulo
Illegal request size.

Transport of request sense fails()
Driver attempted to submit a request sense command and failed.

Transport rejected()
Host adapter driver was unable to accept a command.

Unable to read label
Failure to read disk label.

Unit does not respond to selection
Drive went offline; probably powered down.

NOTES DVD-ROM media containing DVD-Video data may follow/adhere to the requirements of content scrambling system or copy protection scheme. Reading of copy-protected sector will cause I/O error. Users are advised to use the appropriate playback software to view video contents on DVD-ROM media containing DVD-Video data.

NAME	sdt – DTrace statically defined tracing provider
DESCRIPTION	<p>The <code>sdt</code> driver is a DTrace dynamic tracing provider that performs dynamic instrumentation at statically-defined locations in the Solaris kernel.</p> <p>The <code>sdt</code> provider allows kernel developers to explicitly create probes at formally designated locations in the operating system kernel and loadable modules, allowing the implementor to consciously choose the points in their code that are desired probe points, and to convey some semantic knowledge about that point with the choice of probe name and a relevant set of arguments.</p> <p>The <code>sdt</code> driver is not a public interface and you access instrumentation offered by this provider through DTrace. Refer to the <i>Solaris Dynamic Tracing Guide</i> for a description of the public documented interfaces available for the DTrace facility and the probes offered by the <code>sdt</code> provider.</p>
ATTRIBUTES	See <code>attributes(5)</code> for a description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWdtrp
Interface Stability	Private

SEE ALSO `dtrace(1M)`, `attributes(5)`, `dtrace(7D)`
Solaris Dynamic Tracing Guide

se(7D)

NAME	se – Siemens 82532 ESCC serial communications driver
SYNOPSIS	<code>se@bus_address:port_name [,cu]</code>
DESCRIPTION	<p>The <code>se</code> module is a loadable STREAMS driver that provides basic support for the 82532 ESCC hardware and basic asynchronous and synchronous communication support. This manual page describes the asynchronous protocol interface; for information on the synchronous interface, please see the <code>se_hdlc(7D)</code> manual page.</p> <p>Note – This module is affected by the setting of specific eeprom variables. For information on parameters that are persistent across reboots, see the <code>eeprom(1M)</code> man page.</p> <p>The platform specific device bus address for the <code>se</code> module is <code>bus_address</code>. The <code>se</code> module's <code>port_name</code> is a single letter (a-z).</p> <p>Note – During boot up, <code>ttya/b</code> characteristics are read from the <code>/kernel/drv/options.conf</code> file and changed from the PROM defaults to reflect Solaris defaults. Messages displayed on the console after this point are based on settings in that file. If you switch a characteristic, (for example, the baud rate of the console terminal), you must revise the <code>/kernel/drv/options.conf</code> or the console will be configured to an unusable configuration and console messages will be garbled by the mismatched serial port settings.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The Siemens 82532 provides two serial input/output channels capable of supporting a variety of communication protocols. A typical system will use one of these devices to implement two serial ports (<code>port_name</code>), usually configured for RS-423 (which also supports most RS-232 equipment). The Siemens 82532 uses 64 character input and output FIFOs to reduce system overhead. When receiving characters, the CPU is notified when 32 characters have arrived (one-half of receive buffer is full) or no character has arrived in the time it would take to receive four characters at the current baud rate.</p> <p>When sending characters, the Siemens 82532 places the first 64 characters to be sent into its output FIFO and then notifies the CPU when it is half empty (32 characters left). Because the <code>se</code> module waits for the Siemens 82532 to transmit the remaining characters within its output FIFO before making requested changes, delays may occur when the port's attributes are being modified.</p> <p>The <code>se</code> module implements CTS/RTS flow control in hardware. To prevent data overruns, remove CTS/RTS flow control responsibility from the CPU during periods of high system load.</p> <p>In <code>async</code> mode (obtained by opening <code>/dev/cua/[a-z]</code>, <code>/dev/term/[a-z]</code> or <code>/dev/tty[a-z]</code>), the driver supports the <code>termio(7I)</code> device control functions specified by flags in the <code>c_cflag</code> word of the <code>termios</code> structure, and by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code>, or <code>INPCK</code> flags in the <code>c_iflag</code> word. All other <code>termio(7I)</code> functions must be performed by STREAMS modules pushed atop the driver. When a device is opened, the <code>ldterm(7M)</code> and <code>ttcompat(7M)</code> STREAMS modules are automatically pushed on top of the stream, providing the standard <code>termio</code> interface.</p>

Each of the following are valid name space entries: `/dev/cua/[a-z]`, `/dev/term/[a-z]`, and `/dev/tty[a-z]`. The number of entries used in this name space are machine dependent. The `/dev/tty[a-z]` device names exist only if the *SunOS 4.x Binary Compatibility Package* is installed. The `/dev/tty[a-z]` device names are created by the `ucblinks` command, which is available only with the *SunOS 4.x Binary Compatibility Package*.

You can connect a single tty line to a modem for incoming and outgoing calls using a special feature controlled by the minor device number. By accessing character-special devices with names of the form `/dev/cua/[a-z]`, it is possible to open a port without the Carrier Detect signal being asserted, either through hardware or an equivalent software mechanism. These devices are commonly known as dial-out lines.

After a `/dev/cua/[a-z]` line is opened, the corresponding tty line cannot be opened until the `/dev/cua/[a-z]` line is closed. A blocking open will wait until the `/dev/cua/[a-z]` line is closed (which will drop Data Terminal Ready and Carrier Detect) and carrier is detected again. A non-blocking open will return an error. If the tty line has been opened successfully (usually only when carrier is recognized on the modem), the corresponding `/dev/cua/[a-z]` line cannot be opened. This allows a modem to be attached to a device, (for example, `/dev/term/[a-z]` renamed from `/dev/tty[a-z]`) and used for dial-in (by enabling the line for login in `/etc/inittab`) and dial-out (by `tip(1)` or `uucp(1C)`) as `/dev/cua/[a-z]` when no one is logged in on the line.

IOCTLS The `se` module supports the standard set of `termio ioctl()` calls.

Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls.

The state of the DCD, CTS, RTS, and DTR interface signals can be queried through the use of the `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` arguments to the `TIOCMGET ioctl` command, respectively. Due to hardware limitations, only the RTS and DTR signals may be set through their respective arguments to the `TIOCMSET`, `TIOCMBS`, and `TIOCMBSIC ioctl` commands.

The input and output line speeds may be set to all baud rates supported by `termio`. Input and output line speeds cannot be set independently; when you set the output speed, the input speed is automatically set to the same speed.

When using baud rates over 100,000 baud, the software changes the line driver configuration to handle the higher data rates. This action decreases the theoretical maximum cable length from 70 meters to 30 meters.

When the `se` module is used to service the serial console port, it supports a BREAK condition that allows the system to enter the debugger or the monitor. The BREAK condition is generated by hardware and it is usually enabled by default. A BREAK condition originating from erroneous electrical signals cannot be distinguished from one deliberately sent by remote DCE. Due to a risk of incorrect sequence interpretation, SLIP and certain other binary protocols should not be run over the

se(7D)

serial console port when Alternate Break sequence is in effect. Although PPP is a binary protocol, it is able to avoid these sequences using the ACCM feature in *RFC 1662*. For Solaris PPP 4.0, you do this by adding the following line to the `/etc/ppp/options` file (or other configuration files used for the connection; see `pppd(1M)` for details):

```
asynmap 0x00002000
```

By default, the Alternate Break sequence is a three character sequence: carriage return, tilde and control-B (CR ~ CTRL-B), but may be changed by the driver. For information on breaking (entering the debugger or monitor), see `kadb(1M)` and `kb(7M)`.

ERRORS An `open()` will fail under the following conditions:

- ENXIO The unit being opened does not exist.
- EBUSY The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open.
- EBUSY The port is in use by another serial protocol.
- EBUSY The unit has been marked as exclusive-use by another process with a `TIOCEXCL ioctl()` call.
- EINTR The open was interrupted by the delivery of a signal.

- FILES**
- `/dev/cua/[a-z]` dial-out tty lines
 - `/dev/term/[a-z]` dial-in tty lines
 - `/dev/tty[a-z]` binary compatibility package device names
 - `/dev/se_hdlc[0-9]` synchronous devices - see `se_hdlc(7D)`.
 - `/dev/se_hdlc` synchronous control clone device
 - `/kernel/drv/options.conf` system wide default device driver properties

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO `tip(1)`, `kadb(1M)`, `ucblinks(1B)`, `cu(1C)`, `uucp(1C)`, `eeprom(1M)`, `ports(1M)`, `pppd(1M)`, `ioctl(2)`, `open(2)`, `attributes(5)`, `zs(7D)`, `zsh(7D)`, `se_hdlc(7D)`, `termio(7I)`, `ldterm(7M)`, `ttcompat(7M)`, `kb(7M)`

DIAGNOSTICS `sen : fifo overrun` The Siemens 82532 internal FIFO received more data than it could handle. This indicates that Solaris was not servicing data interrupts fast enough and suggests a system with too many interrupts or a data line with a data rate that is too high.

se(7D)

`sen` : buffer overrun

The `se` module was unable to store data it removed from the Siemens 82532 FIFO. The user process is not reading data fast enough, and suggests an overloaded system. If possible, the application should enable flow control (either CTSRTS or XONXOFF) to allow the driver to backpressure the remote system when the local buffers fill up.

se_hdlc(7D)

NAME	se_hdlc – on-board high-performance serial HDLC interface
SYNOPSIS	<i>se@bus_address:port_number</i> [, <i>hdlc</i>]
DESCRIPTION	<p>The <i>se_hdlc</i> devices are a synchronous hdlc-framing interface for the <i>se</i> serial devices. Both built-in serial ports (<i>port_number</i>) on platforms which have the <i>se</i> serial devices, support synchronous data transfer at a maximum rate of 384 kbps. <i>bus_address</i> is the platform specific <i>se</i> device bus address. <i>port_number</i> is a single digit number (0-9).</p>
APPLICATION PROGRAMMING INTERFACE	<p>The <i>se_hdlcn</i> devices provide a data path which supports the transfer of data via <i>read(2)</i> and <i>write(2)</i> system calls, as well as <i>ioctl(2)</i> calls. Data path opens are exclusive in order to protect against injection or diversion of data by another process.</p> <p>The <i>se_hdlc</i> device provides a separate control path for use by programs that need to configure or monitor a connection independent of any exclusive access restrictions imposed by data path opens. Up to three control paths may be active on a particular serial channel at any one time. Control path accesses are restricted to <i>ioctl(2)</i> calls only; no data transfer is possible.</p> <p>When used in synchronous modes, the SAB 82532 ESCC supports several options for clock sourcing and data encoding. Both the transmit and receive clock sources can be set to be the external Transmit clock (TRxC), external Receive Clock (RTxC), the internal Baud Rate Generator (BRG), or the output of the ESCC 's Digital Phase-Lock Loop (DPLL).</p> <p>The BRG is a programmable divisor that derives a clock frequency from the PCLK input signal to the ESCC. The programmed baud rate is translated into a floating point (6-bit mantissa, 4-bit exponent) number time constant that is stored in the ESCC.</p> <p>A local loopback mode is available, primarily for use by <i>syncloop(1M)</i> for testing purposes, and should not be confused with SDLC loop mode, which is not supported on this interface. Also, an auto-echo feature may be selected that causes all incoming data to be routed to the transmit data line, allowing the port to act as the remote end of a digital loop. Neither of these options should be selected casually, or left in use when not needed.</p> <p>The <i>se</i> driver keeps running totals of various hardware generated events for each channel. These include numbers of packets and characters sent and received, abort conditions detected by the receiver, receive CRC errors, transmit underruns, receive overruns, input errors and output errors, and message block allocation failures. Input errors are logged whenever an incoming message must be discarded, such as when an abort or CRC error is detected, a receive overrun occurs, or when no message block is available to store incoming data. Output errors are logged when the data must be discarded due to underruns, CTS drops during transmission, CTS timeouts, or excessive watchdog timeouts caused by a cable break.</p>
IOCTLS	The <i>se</i> driver supports the following <i>ioctl()</i> commands.

S_IOCGETMODE	Return a struct <code>scc_mode</code> containing parameters currently in use. These include the transmit and receive clock sources, boolean loopback and NRZI mode flags and the integer baud rate.
S_IOCSETMODE	The argument is a struct <code>scc_mode</code> from which the ESCC channel will be programmed.
S_IOCGETSTATS	Return a struct <code>sl_stats</code> containing the current totals of hardware-generated events. These include numbers of packets and characters sent and received by the driver, aborts and CRC errors detected, transmit underruns, and receive overruns.
S_IOCCLRSTATS	Clear the hardware statistics for this channel.
S_IOCGETSPEED	Returns the currently set baud rate as an integer. This may not reflect the actual data transfer rate if external clocks are used.
S_IOCGETMCTL	Returns the current state of the CTS and DCD incoming modem interface signals as an integer.

The following structures are used with `se_hdlc ioctl()` commands:

```

struct scc_mode {
    char sm_txclock;    /* transmit clock sources */
    char sm_rxclock;    /* receive clock sources */
    char sm_iflags;     /* data and clock inversion flags (non-zsh) */
    uchar_t sm_config; /* boolean configuration options */
    int sm_baudrate;    /* real baud rate */
    int sm_retval;      /* reason codes for ioctl failures */
};

struct sl_stats {
    long ipack;        /* input packets */
    long opack;        /* output packets */
    long ichar;        /* input bytes */
    long ochar;        /* output bytes */
    long abort;        /* abort received */
    long crc;          /* CRC error */
    long cts;          /* CTS timeouts */
    long dcd;          /* Carrier drops */
    long overrun;      /* receive overrun */
    long underrun;     /* transmit underrun */
    long ierror;       /* input error */
    long oerror;       /* output error */
    long nobuffers;    /* receive side memory allocation failure */
};

```

ERRORS An `open()` will fail if a STREAMS message block cannot be allocated or under the following conditions:

ENXIO The unit being opened does not exist.

se_hdlc(7D)

EBUSY The device is in use by another serial protocol.

An `ioctl()` will fail under the following conditions:

EINVAL An attempt was made to select an invalid clocking source.

EINVAL The baud rate specified for use with the baud rate generator would translate to a null time constant in the ESCC's registers.

FILES `/dev/se_hdlc[0-1]`, character-special devices
`/dev/se_hdlc`
`/usr/include/sys/ser_sync.h` header file specifying synchronous serial communication definitions

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO `syncinit(1M)`, `syncloop(1M)`, `syncstat(1M)`, `ioctl(2)`, `open(2)`, `read(2)`, `write(2)`, `attributes(5)`, `se(7D)`, `zsh(7D)`

Siemens ESCC2 SAB 82532 Enhanced Serial Communication Controller User's Manual

DIAGNOSTICS `se_hdlc clone open failed, no memory, rq=nnn`
A kernel memory allocation failed for one of the private data structures. The value of `nnn` is the address of the read queue passed to `open(2)`.

`se_hdlc: clone device must be attached before use!`
An operation was attempted through a control path before that path had been attached to a particular serial channel.

`se_hdlcn: not initialized, can't send message`
An `M_DATA` message was passed to the driver for a channel that had not been programmed at least once since the driver was loaded. The ESCC's registers were in an unknown state. The `S_IOCSETMODE` `ioctl` command performs the programming operation.

`se_hdlc_start: Invalid message type d on write queue`
driver received an invalid message type from streams.

`se_hdlcn: transmit hung`
The transmitter was not successfully restarted after the watchdog timer expired. This is usually caused by a bad or disconnected cable.

NAME	ses – SCSI enclosure services device driver
SYNOPSIS	<code>ses@target, lun</code>
DESCRIPTION	<p>The <code>ses</code> device driver is an interface to SCSI enclosure services devices. These devices sense and monitor the physical conditions within an enclosure as well as allow access to the status reporting and configuration features of the enclosure (such as indicator LEDs on the enclosure.)</p> <p><code>ioctl(9E)</code> calls may be issued to <code>ses</code> to determine the state of the enclosure and to set parameters on the enclosure services device.</p> <p>No <code>ses</code> driver properties are defined. Use the <code>ses.conf</code> file to configure the <code>ses</code> driver.</p>
EXAMPLES	<p>EXAMPLE 1 <code>ses.conf</code> File Format</p> <p>The following is an example of the <code>ses.conf</code> file format:</p> <pre> # # Copyright (c) 1996, by Sun Microsystems, Inc. # All rights reserved. # # #ident "@(#)ses.conf 1.1 97/02/10 SMI" # name="ses" parent="sf" target=15; name="ses" parent="SUNW,pln" port=0 target=15; name="ses" parent="SUNW,pln" port=1 target=15; name="ses" parent="SUNW,pln" port=2 target=15; name="ses" parent="SUNW,pln" port=3 target=15; name="ses" parent="SUNW,pln" port=4 target=15; name="ses" parent="SUNW,pln" port=5 target=15; name="ses" class="scsi" target=15 lun=0; </pre>
IOCTLS	<p>The SES driver currently supports the SES, SAFTE and SEN enclosure service chipsets. SEN and SAFTE protocols are translated internally in the driver into SES compliant data structures. This enables the SES driver to work seamlessly with different protocols and eliminates the need to enhance user applications.</p> <p>SESIOC_GETNOBJ Returns an unsigned integer that represents the number of SES data structures in the enclosure services chip.</p> <p>SESIOC_GETOBJMAP Returns a size array containing <code>ses_object</code> elements communicated through <code>SESIOC_GETNOBJ()</code>. <code>ses_object</code> is defined in <code>sesio.h</code>.</p> <p>SESIOC_INIT Instructs the device to perform a self-diagnostic test. Currently SES & SEN devices always return success.</p>

ses(7D)

SESIOC_GETENCSTAT

Returns an unsigned character that represents status enclosure as defined by Table 25 in Section 7.1.2 of the SES specification *NCITS 305-199x*.

SESIOC_GETOBJSTAT

This ioctl is passed an `ses_objarg` containing the `obj_id` you want to set, then fills in the remaining fields according to element status page of the SES specification.

SESIOC_SETOBJSTAT

Sets options in the control field. You set control field options by filling out all fields in `ses_objarg`. Field definitions are presented in Section 7.2.2 of the SES specification.

FILES /kernel/drv/ses.conf Driver configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO `driver.conf(4)`, `attributes(5)`, `esp(7D)`, `isp(7D)`, `ioctl(9E)`

NAME	sesio – enclosure services device driver interface				
SYNOPSIS	<code>#include <sys/sesio.h></code>				
DESCRIPTION	The <i>ses</i> device driver provides the following ioctls as a means to access SCSI enclosure services devices.				
IOCTLS	<p>The <i>ses</i> driver supports the following ioctls:</p> <p><code>SES_IOCTL_GETSTATE</code> This ioctl obtains enclosure state in the <code>ses_ioctl</code> structure.</p> <p><code>SES_IOCTL_SETSTATE</code> This ioctl is used to set parameters on the enclosure services device. The <code>ses_ioctl</code> structure is used to pass information into the driver.</p>				
ERRORS	<p><code>EIO</code> The <i>ses</i> driver was unable to obtain data from the enclosure services device or the data transfer could not be completed.</p> <p><code>ENOTTY</code> The <i>ses</i> driver does not support the requested ioctl function.</p> <p><code>ENXIO</code> The enclosure services device does not exist.</p> <p><code>EFAULT</code> The user specified a bad data length.</p>				
STRUCTURES	<p>The <code>ses_ioctl</code> structure has the following fields:</p> <pre>uint32_t; /* Size of buffer that follows */ uint8_t page_code: /* Page to be read/written */ uint8_t reserved[3]; /* Reserved; Set to 0 */ uint8_t buffer[1]; /* Size arbitrary, user specifies */</pre>				
EXAMPLES	<p>EXAMPLE 1 Using the <code>SES_IOCTL_GETSTATE</code> ioctl</p> <p>The following example uses the <code>SES_IOCTL_GETSTATE</code> ioctl to recover 20 bytes of page 4 from a previously opened device.</p> <pre>char abuf[30]; struct ses_ioctl *sesp; int status; sesp = (ses_ioctl *)abuf; sesp->size = 20; sesp->page_code = 4; status = ioctl(fd, SES_IOCTL_GETSTATE, abuf);</pre>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">Architecture</td> <td style="text-align: center;">SPARC</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	SPARC				
SEE ALSO	<code>ses(7D)</code> , <code>ioctl(9E)</code>				

sf(7D)

NAME	sf – SOC+ FC-AL FCP Driver				
SYNOPSIS	<code>sf@port, 0</code>				
DESCRIPTION	<p>The <code>sf</code> driver is a SCSI compliant nexus driver which supports the Fibre Channel Protocol for SCSI on Private Fibre Channel Arbitrated loops. An SBus card called the SOC+ card (see socal(7D)) connects the Fibre Channel loop to the host system.</p> <p>The <code>sf</code> driver interfaces with the SOC+ device driver, socal(7D), the SCSI disk target driver, ssd(7D), and the SCSI-3 Enclosure Services driver, ses(7D). It only supports SCSI devices of type disk and ses.</p> <p>The <code>sf</code> driver supports the standard functions provided by the SCSI interface. The driver supports auto request sense and tagged queueing by default.</p> <p>The driver requires that all devices have unique hard addresses defined by switch settings in hardware. Devices with conflicting hard addresses will not be accessible.</p>				
FILES	<table><tr><td><code>/platform/architecture/kernel/drv/sf</code></td><td>ELF kernel module</td></tr><tr><td><code>/platform/architecture/kernel/drv/sf.conf</code></td><td><code>sf</code> driver configuration file</td></tr></table>	<code>/platform/architecture/kernel/drv/sf</code>	ELF kernel module	<code>/platform/architecture/kernel/drv/sf.conf</code>	<code>sf</code> driver configuration file
<code>/platform/architecture/kernel/drv/sf</code>	ELF kernel module				
<code>/platform/architecture/kernel/drv/sf.conf</code>	<code>sf</code> driver configuration file				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>SPARC</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Architecture	SPARC
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Architecture	SPARC				
SEE ALSO	<code>luxadm(1M)</code> , <code>prtconf(1M)</code> , <code>driver.conf(4)</code> , socal(7D) , ssd(7D) <i>Writing Device Drivers</i> <i>ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL)</i> <i>ANSI X3.269-1996, Fibre Channel Protocol for SCSI (FCP)</i> <i>ANSI X3.270-1996, SCSI-3 Architecture Model (SAM)</i> <i>Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)</i>				
DIAGNOSTICS	<p>In addition to being logged, the messages below may display on the system console.</p> <p>The first set of messages indicate that the attachment was unsuccessful, and will only display while the <code>sf</code> driver is initially attempting to attach. Each message is preceded by <code>sf%d</code>, where <code>%d</code> is the instance number of the <code>sf</code> device.</p> <p>Failed to alloc soft state Driver was unable to allocate space for the internal state structure. Driver did not attach to device, SCSI devices will be inaccessible.</p> <p>Bad soft state Driver requested an invalid internal state structure. Driver did not attach to device, SCSI devices will be inaccessible.</p>				

Failed to obtain transport handle

Driver was unable to obtain a transport handle to communicate with the social driver. Driver did not attach to device, SCSI devices will be inaccessible

Failed to allocate command/response pool

Driver was unable to allocate space for commands and responses. Driver did not attach to device, SCSI devices will be inaccessible.

Failed to allocate kmem cache

Driver was unable to allocate space for the packet cache. Driver did not attach to device, SCSI devices will be inaccessible.

Failed to allocate dma handle for

Driver was unable to allocate a dma handle for the loop map. Driver did not attach to device, SCSI devices will be inaccessible.

Failed to allocate lilp map

Driver was unable to allocate space for the loop map. Driver did not attach to device, SCSI devices will be inaccessible.

Failed to bind dma handle for

Driver was unable to bind a dma handle for the loop map. Driver did not attach to device, SCSI devices will be inaccessible.

Failed to attach

Driver was unable to attach for some reason that may be printed. Driver did not attach to device, SCSI devices will be inaccessible.

The next set of messages may display at any time. The full device pathname, followed by the shorter form described above, will precede the message.

Invalid lilp map

The driver did not obtain a valid lilp map from the social driver. SCSI device will be inaccessible.

Target t, AL-PA x and hard

The device with a switch setting t has an AL-PA x which does not match its hard address y. The device will not be accessible.

Duplicate switch settings

The driver detected devices with the same switch setting. All such devices will be inaccessible.

WWN changed on target t

The World Wide Name (WWN) has changed on the device with switch setting t.

Target t, unknown device type

The driver does not know the device type reported by the device with switch setting t.

sgen(7D)

NAME	sgen – Generic SCSI device driver
SYNOPSIS	<pre>#include <sys/scsi/targets/sgendef.h> sgen@target, lun:<devtype></pre>
DESCRIPTION	<p>The <code>sgen</code> driver exports the <code>uscsci(7I)</code> interfaces to user processes. The <code>sgen</code> driver can be configured to bind to SCSI devices for which no system driver is available. Examples of such devices include SCSI scanners and SCSI processor devices.</p>
SECURITY AND DATA INTEGRITY	<p>Typically, drivers which export the <code>uscsci(7I)</code> interface unconditionally require that the user present superuser credentials. The <code>sgen</code> driver does not, and relies on the filesystem permissions on its device special file to govern who may access that device. By default, access is restricted and device nodes created by the <code>sgen</code> driver are readable and writable by the superuser exclusively.</p> <p>It is important to understand that SCSI devices coexisting on the same SCSI bus may potentially interact with each other. This may result from firmware bugs in SCSI devices, or may be made to happen programmatically by sending appropriate SCSI commands to a device. Potentially, any application controlling a device via the <code>sgen</code> driver can introduce data integrity or security problems in that device or any other device sharing the same SCSI bus.</p> <p>Granting unprivileged users access to an <code>sgen</code>-controlled SCSI device may create other problems. It may be possible for a user to instruct a target device to gather data from another target device on the same bus. It may also be possible for malicious users to install new firmware onto a device to which they are granted access. In environments where security is a concern but user access to devices controlled by the <code>sgen</code> driver is nonetheless desired, it is recommended that the devices be separated onto a dedicated SCSI bus to mitigate the risk of data corruption and security violations.</p>
CONFIGURATION	<p>The <code>sgen</code> driver is configurable via the <code>sgen.conf</code> file. In addition to standard SCSI device configuration directives (see <code>scsi(4)</code>), administrators can set several additional properties for the <code>sgen</code> driver.</p> <p>By default, the <code>sgen</code> driver will not claim or bind to any devices on the system. To do so, it must be configured by the administrator using the <code>inquiry-config-list</code> and/or the <code>device-type-config-list</code> properties.</p> <p>As with other SCSI drivers, the <code>sgen.conf</code> configuration file enumerates the targets <code>sgen</code> should use. See <code>scsi(4)</code> for more details. For each target enumerated in the <code>sgen.conf</code> file, the <code>sgen</code> driver sends a SCSI INQUIRY command to gather information about the device present at that target. The <code>inquiry-config-list</code> property specifies that the <code>sgen</code> driver should bind to a particular device returning a particular set of inquiry data. The <code>device-type-config-list</code> specifies that the <code>sgen</code> driver should bind to every device that is of a particular SCSI device type. When examining the device, the <code>sgen</code> driver tests to see if it matches an entry in the <code>device-type-config-list</code> or the <code>inquiry-config-list</code>. For more detail on these two properties, see the PROPERTIES section.</p>

When a match against the `INQUIRY` data presented by a device is made, the `sgen` driver attaches to that device and creates a device node and link in the `/devices` and `/dev` hierarchies. See the `FILES` section for more information about how these files are named.

It is important for the administrator to ensure that devices claimed by the `sgen` driver do not conflict with existing target drivers on the system. For example, if the `sgen` driver is configured to bind to a direct access device, the standard `sd.conf` file will usually cause `sd` to claim the device as well. This can cause unpredictable results. In general, the `uscsi(7I)` interface exported by `sd(7D)` or `st(7D)` should be used to gain access to direct access and sequential devices.

The `sgen` driver is disabled by default. The `sgen.conf` file is shipped with all of the `'name="sgen" class="scsi" target=...'` entries commented out to shorten boot time and to prevent the driver from consuming kernel resources. To use the `sgen` driver effectively on desktop systems, simply uncomment all of the `name="sgen"` lines in `sgen.conf` file. On larger systems with many SCSI controllers, carefully edit the `sgen.conf` file so that `sgen` binds only where needed. Refer to `driver.conf(4)` for further details.

PROPERTIES

`inquiry-config-list`

The `inquiry-config-list` property is a list of pairs of strings that enumerates a list of specific devices to which the `sgen` driver will bind. Each pair of strings is referred to as `<vendorid, productid>` in the discussion below.

`vendorid`

is used to match the Vendor ID reported by the device. The SCSI specification limits Vendor IDs to eight characters. Correspondingly, the length of this string should not exceed eight characters. As a special case, `"*"` may be used as a wildcard which matches any Vendor ID. This is useful in situations where more than one vendor produces a particular model of a product. `vendorid` is matched against the Vendor ID reported by the device in a case-insensitive manner.

`productid`

is used to match the product ID reported by the device. The SCSI specification limits product IDs to sixteen characters (unused characters are filled with the whitespace characters). Correspondingly, the length of `productid` should not exceed sixteen characters. When examining the product ID of the device, `sgen` examines the length `l` of `productid` and performs a match against only the first `l` characters in the device's product ID. `productid` is matched against the product ID reported by the device in a case-insensitive manner.

For example, to match some fictitious devices from ACME corp, the `inquiry-config-list` can be configured as follows:

```
inquiry-config-list =      "ACME",          "UltraToast 3000",
                          "ACME",          "UltraToast 4000",
```

sgen(7D)

```
"ACME", "UltraToast 5000";
```

To match "UltraToast 4000" devices, regardless of vendor, `inquiry-config-list` is modified as follows:

```
inquiry-config-list = "*", "UltraToast 4000";
```

To match every device from ACME in the "UltraToast" series (i.e UltraToast 3000, 4000, 5000, ...), `inquiry-config-list` is modified as follows:

```
inquiry-config-list = "ACME" "UltraToast";
```

Whitespace characters *are* significant when specifying `productId`. For example, a `productId` of "UltraToast 1000" is fifteen characters in length. If a device reported its ID as "UltraToast 10000", the `sgen` driver would bind to it because only the first fifteen characters are considered significant when matching. To remedy this situation, specify `productId` as "UltraToast 1000 ", (note trailing space). This forces the `sgen` driver to consider all sixteen characters in the product ID to be significant.

`device-type-config-list`

The `device-type-config-list` property is a list of strings that enumerate a list of device types to which the `sgen` driver will bind. The valid device types correspond to those defined by the *SCSI-3 SPC Draft Standard, Rev. 11a*. These types are:

Type Name	Inquiry Type ID
direct	0x00
sequential	0x01
printer	0x02
processor	0x03
worm	0x04
rodirect	0x05
scanner	0x06
optical	0x07
changer	0x08
comm	0x09
prepress1	0x0a

Type Name	Inquiry Type ID
prepress2	0x0b
array_ctrl	0x0c
ses	0x0d
rbc	0x0e
ocrw	0x0f
bridge	0x10
type_unknown	0x1f

Alternately, you can specify device types by INQUIRY type ID. To do this, specify `type_0x<typenum>` in the `sgen-config-list`. Case is not significant when specifying device type names.

`sgen-diag`

The `sgen-diag` property sets the diagnostic output level. This property can be set globally and/or per target/lun pair. `sgen-diag` is an integer property, and can be set to 0, 1, 2 or 3. Illegal values will silently default to 0. The meaning of each diagnostic level is as follows:

- 0 No error reporting [default]
- 1 Report driver configuration information, unusual conditions, and indicate when sense data has been returned from the device.
- 2 Trace the entry into and exit from routines inside the driver, and provide extended diagnostic data. No error reporting [default].
- 3 Provide detailed output about command characteristics, driver state, and the contents of each CDB passed to the driver.

In ascending order, each level includes the diagnostics that the previous level reports. See the IOCTLS section for more information on the `SGEN_IOC_DIAG` ioctl.

FILES

`sgen.conf`

Driver configuration file. See CONFIGURATION for more details.

`/dev/scsi/<devtype>/cntndn`

The `sgen` driver categorizes each device in a separate directory by its SCSI device type. The files inside the directory are named according to their controller number, target ID and LUN as follows:

`cn` is the controller number, `tn` is the SCSI target id and `dn` is the SCSI LUN

This is analogous to the `{controller;target;device}` naming scheme, and the controller numbers correspond to the same controller numbers which are used for naming disks. For example, `/dev/dsk/c0t0d0s0` and `/dev/scsi/scanner/c0t5d0` are both connected to controller `c0`.

sgen(7D)

IOCTLS	<p>The <code>sgen</code> driver exports the uscsi(7I) interface for each device it manages. This allows a user process to talk directly to a SCSI device for which there is no other driver installed in the system. Additionally, the <code>sgen</code> driver supports the following ioctls:</p> <p><code>SGEN_IOC_READY</code> Send a <code>TEST UNIT READY</code> command to the device and return 0 upon success, non-zero upon failure. This ioctl accepts no arguments.</p> <p><code>SGEN_IOC_DIAG</code> Change the level of diagnostic reporting provided by the driver. This ioctl accepts a single integer argument between 0 and 3. The levels have the same meaning as in the <code>sgen-diag</code> property discussed in <code>PROPERTIES</code> above.</p>
ERRORS	<p><code>EBUSY</code> The device was opened by another thread or process. The driver maintains a strict exclusive-open policy for each device.</p> <p><code>ENXIO</code> During opening, the device did not respond to a <code>TEST UNIT READY</code> SCSI command.</p> <p><code>ENOTTY</code> Indicates that the device does not support the requested ioctl function.</p>
EXAMPLES	<p>Here is an example of how <code>sgen</code> can be configured to bind to scanner devices on the system:</p> <pre>device-type-config-list = "scanner";</pre> <p>The administrator should subsequently uncomment the appropriate <code>name="sgen" . . .</code> lines for the SCSI target ID to which the scanner corresponds. In this example, the scanner is at target 4.</p> <pre>name= "sgen" class= "scsi" target=4 lun=0;</pre> <p>If it is expected that the scanner will be moved from target to target over time, or that more scanners might be added in the future, it is recommended that all of the <code>name="sgen" . . .</code> lines be uncommented, so that <code>sgen</code> checks all of the targets on the bus.</p> <p>For large systems where boot times are a concern, it is recommended that the <code>parent=""</code> property be used to specify which SCSI bus <code>sgen</code> should examine.</p>
SEE ALSO	<p>driver.conf(4), scsi(4), sd(7D), st(7D), uscsi(7I)</p> <p><i>Writing Device Drivers</i></p> <p><i>ANSI Small Computer System Interface-2 (SCSI-2)</i></p> <p><i>SCSI-3 SPC Draft Standard, Rev. 11a</i></p>

NAME	sip – SIP Proxy/registrar/redirect server
DESCRIPTION	Solaris supports deployment of VoIP/SIP services by providing an <i>RFC 3261</i> –compliant SIP proxy/registrar/redirect server called SER from <i>iptel.org</i> . See the <code>ser(8)</code> man page under <code>/usr/sfw/man</code> .
FILES	<code>/etc/sfw/ser/ser.cfg</code> <code>/etc/sfw/ser/README.solaris.ser</code> <code>/usr/sfw/share/doc/ser/README</code>

sk98sol(7D)

NAME	sk98sol – SysKonnnect Gigabit Ethernet SK-98xx device driver
SYNOPSIS	<pre>/dev/skge /kernel/drv/sk98sol</pre>
DESCRIPTION	The <code>sk98sol</code> driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface (DLPI), over a SysKonnnect Gigabit Ethernet adapter (SK-98xx series). The driver supports multiple installed SysKonnnect SK-98xx adapters. Functions include chip initialization, frame transmit and receive, multicast and promiscuous support, and error recovery and reporting.
APPLICATION PROGRAMMING INTERFACE	The driver provides the <code>/dev/skge</code> cloning character-special device as well as per-adapter character-special devices <code>/dev/skgex</code> , where x represents the device instance number.
sk98sol and DLPI	<p>The <code>sk98sol</code> driver is a Style 1 and Style 2 Data Link Service (DLS) provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code> type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <code><sys/dlpi.h></code>. See <code>dlpi(7P)</code>.</p> <p>An explicit <code>DL_ATTACH_REQ</code> message by the user is required to associate the opened stream with a particular device (<i>ppa</i>). This is unnecessary and invalid for DLPI Style 1. The <i>ppa</i> ID is interpreted as an <code>unsigned long</code> data type and indicates the corresponding device instance (unit) number. An error (<code>DL_ERROR_ACK</code>) is returned by the driver if the <i>ppa</i> field value does not correspond to a valid device instance number for the system.</p> <p>The device is initialized on first attach and de-initialized (stopped) upon last detach. Valid device numbers for all detected adapters are displayed on the console at driver startup time and are written to the <code>/var/adm/messages</code> log file.</p> <p>The values returned in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> request are:</p> <ul style="list-style-type: none">■ Maximum SDU is 1500 (9000 if <i>JumboFrames</i> are enabled).■ Minimum SDU is 0.■ DLSAP address length is 8 bytes.■ MAC type is <code>DL_CSMACD</code>.■ SAP length value is -2, meaning the physical address component is followed immediately by a 2-byte SAP component within the DLSAP address.■ Service mode is <code>DL_CLDLS</code>.■ Optional quality of service (QOS) support is not included; as a result, the QOS field values are 0.■ Provider style is <code>DL_STYLE2</code>.■ Version is <code>DL_VERSION_2</code>.

OPTIONS Parameters are set in the `/kernel/drv/sk98sol.conf` configuration file, which is created during installation. See `driver.conf(4)`. You can edit the `/kernel/drv/sk98sol.conf` file to reflect your settings and reboot the system to use the new parameter values. If the file exists prior to driver installation, the new parameter values will be used as soon as the driver is installed.

String parameter values must be surrounded with double quotes ("`\"`"), while integer parameter values are not. Parameter names and values are case sensitive and you should use them exactly as shown.

Note – You can increase `sk98sol` performance by tuning certain TCP and UDP parameters. However, you should be aware that this act may adversely impact the performance of other network cards.

To tune specific parameters to increase `sk98sol` performance, do the following:

```

ndd -set /dev/tcp tcp_xmit_hiwat      65536
ndd -set /dev/tcp tcp_xmit_lowat     32768
ndd -set /dev/tcp tcp_rcv_hiwat     65536

ndd -set /dev/udp udp_xmit_hiwat     65536
ndd -set /dev/udp udp_xmit_lowat     327
ndd -set /dev/udp udp_rcv_hiwat     65536

```

Per-Port Parameters

The parameters discussed in this section can be set for each port on the adapter.

In each of the following descriptions, ? represents port A or B.

AutoNegotiation_?

Type:	String
Values:	On, Off, Sense
Default:	Sense (for SK-984x fiber adapters).
Default:	On (for SK-982x copper adapters.)

The Sense value automatically detects whether the link partner supports autonegotiation. If your link partner is configured to half duplex with autonegotiation turned off, set the *AutoNegotiation_?* and *DuplexCapabilities_?* parameters manually. Do *not* set the *AutoNegotiation_?* parameter value to Sense, as it will fail.

Do not use Sense for 1000Base-T (copper) adapters. If Sense is selected, it will be mapped to On automatically.

DuplexCapabilities_?

Type:	String
-------	--------

Values:	Half, Full, Both
Default:	Both

Set the *DuplexCapabilities_?* parameter only if the *AutoNegotiation_?* parameter is set to the On or Off values. If *AutoNegotiation_?* is set to On, all three *AutoNegotiation_?* values are possible; however, if set to Off, only the Full and Half values are allowed.

Use the *DuplexCapabilities_?* parameter if your link partner does not support all possible combinations.

FlowControl_?

Type:	String
Values:	Sym, SymOrRem, LocSend, None
Default:	SymOrRem

Use the *FlowControl_?* parameter to set the flow control capabilities reported by the port during autonegotiation:

Sym	Symetric flow control, where both link partners are allowed to send PAUSE frames.
SymOrRem	SymetricOrRemote flow control, where both link partners or only the remote partner are allowed to send PAUSE frames.
LocSend	LocalSend flow control, where only the local link partner is allowed to send PAUSE frames.
None	No flow control, where no link partner is allowed to send PAUSE frames.

The *FlowControl_?* parameter is ignored if *AutoNegotiation_?* is set to "Off . "

Role_?

Type:	String
Values:	Auto, Master, Slave
Default:	Auto

Use the *Role_?* parameter only for the SK-9821 and SK-9822 adapters.

**Per-Adapter
Parameters**

1000Base-T communication between two ports requires one port to act as the master (and provide timing information) and the other as slave. Normally, this is negotiated between the two ports during link establishment. If this fails, use the *Role_?* parameter to force the master and slave roles on the ports. If *AutoNegotiation_?* is set to "OFF," then the *Role_?* parameter must be set manually.

PreferredPort

Type:	String
Values:	A, B
Default:	A

Use the *PreferredPort* parameter to force the preferred port to A or B (on two-port NICs). The preferred port is the port selected if both ports are detected as fully functional.

RlmtMode

Type:	integer
Values:	1, 2, 3
Default:	1

RLMT (Redundant Link Management Technology) provides three modes to determine if a port is available for use.

1. Check link state only: use the link state reported by the adapter hardware for each individual port.
2. Check other port: RLMT sends test frames from one port to another and checks if they are received. The ports must be connected to the network that allow LLC test frames to be exchanged (that is, networks without routers between the ports).
3. Check other port and segmentation: RLMT checks the other port and also requests information from the Gigabit Ethernet switch next to each port to determine if the network is segmented between the ports. Only use this mode if you have Gigabit Ethernet switches installed and configured to use the Spanning Tree protocol.

Note that modes 2 and 3 are meant to operate in configurations where a network path exists between the ports on a single adapter. They are *not* designed to work in networks where adapters are connected back-to-back.

JumboFrames

Type:	String
Values:	Off, On
Default:	Off

To enable support for *JumboFrames* (frames with a length of up to 9014 bytes), set `JumboFrames` to "On." Because longer frames reduce operating system overhead, *JumboFrames* increases network throughput.

For full *JumboFrames* support, the maximum transfer unit (MTU) size used by TCP/IP must also be changed by using the `ifconfig(1M)` command. To do this, remove the comment sign (#) before the `ifconfig` line in the `/etc/rcS.d/S50sk98sol` file. You should also ensure that the adapter device number (`skge0`) matches the attach number displayed during system startup. The MTU must be set to 9000, not including the 14 bytes of MAC address header.

JumboFrames can only be used if *all* equipment in your subnetwork supports them; currently many switches do not support *JumboFrames*). Devices without Jumbo Support drop the longer frames (and might report them as error frames). If you experience problems with this, connect two SK-98xx adapters (with *JumboFrames* enabled) back-to-back.)

CopyThreshold

Type:	Integer
Values:	0-1500
Default:	1500

During transmit, the driver relies on the frame's physical memory address to tell the hardware where to find the frame data. Setting up the DMA address can take time on Solaris; it may be more convenient to copy the frame data to a buffer that you have set up in advance. All frames with a length less than or equal to the *CopyThreshold* parameter value are copied into buffers; for longer frames, the real DMA setup is done. By default (without *JumboFrames* support), all frames are copied. You can experiment with this parameter to find out if your system performs better with only smaller frames copied.

To use more complex syntax for setting different parameters on multiple adapters, see `driver.conf(4)`. For example:

```
name="sk98sol" parent="/pci@1f,4000" unit-address="2"
AutoNegotiation_A="Off";
name="sk98sol" parent="/pci@1f,2000" unit-address="2"
AutoNegotiation_B="Sense";
```

DIAGNOSTICS If multiple NICs are installed in the system, the following message may appear on the console and in the `/var/adm/messages` log file:

```
Allocation of descriptor memory failed
```

You can avoid this message by tuning the `lomempages` kernel parameter. By default, the value of this parameter is 36 pages. Each SK-98xx adapter requires a determined number of pages, so increase the value of the `lomempages` parameter in increments of ten pages until all NICs in the system run correctly.

To modify the value of this parameter to 46 pages, append the `set lomempages=46` line to the `/etc/system` file and reboot the system.

FILES	<code>/dev/skge</code>	Character special device
	<code>/dev/skgex</code>	Per-adapter character special device, where <i>x</i> is the adapter <i>ppa</i>
	<code>/kernel/drv/sk98sol</code>	ELF kernel module
	<code>/kernel/drv/sparcv9/sk98sol</code>	64-bit ELF kernel module
	<code>/kernel/drv/sk98sol.conf</code>	Driver configuration file

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	IA, SPARC

SEE ALSO `ifconfig(1M)`, `netstat(1M)`, `driver.conf(4)`, `attributes(5)`, `dlpi(7P)`.

`sk98sol.txt` driver README file — Included in the driver package; also available from www.sysconnect.com.

skfp(7D)

NAME	skfp – SysKconnect FDDI PCI device driver
SYNOPSIS	<code>/dev/skfp</code>
DESCRIPTION	<p>The <code>skfp</code> FDDI driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface (DLPI) over a SysKconnect FDDI PCI adapter. The driver supports multiple installed SysKconnect FDDI PCI adapters. Functions include chip initialization, frame transit and receive, multicast and promiscuous support, and error recovery and reporting.</p> <p>The <code>skfp</code> driver supports all SysKconnect SK-NET FDDI PCI adapters (SK-55xx (32-bit) and SK-58xx (64-bit) series) on 32-bit systems, and the SK-58xx series on 64-bit systems.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The <code>skfp</code> driver provides the <code>/dev/skfp</code> cloning character-special device that accesses all SK-NET FDDI PCI adapters using Data Link Service (DLS) Style 2. It also provides per-adapter character-special devices <code>/dev/skfpx</code>, (where x represents the device instance number) that access a special NIC using DLS Style 1.</p>
skfp and DLPI	<p>The <code>skfp</code> driver is a Style 1 and Style 2 DLS provider. All <code>M_PROTO</code> and <code>M_PCPROTO</code> type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in <code><sys/dlpi.h></code>. See dlpi(7P).</p> <p>An explicit <code>DL_ATTACH_REQ</code> message by the user is required to associate the opened Stream with a particular device (<i>ppa</i>). This is unnecessary and invalid for DLPI Style 1. The <i>ppa</i> ID is interpreted as an unsigned long data type and indicates the corresponding device instance (unit) number. An error (<code>DL_ERROR_ACK</code>) is returned by the driver if the <i>ppa</i> field value does not correspond to a valid device instance number for the system.</p> <p>The device is initialized on first attach and de-initialized (stopped) upon last detach. Valid device numbers for all detected adapters are displayed on the console at driver startup time and written to the <code>/var/adm/messages</code> log file.</p> <p>The values returned in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> request are:</p> <ul style="list-style-type: none">■ Maximum SDU is 4470.■ Minimum SDU is 0.■ DSLAP address length is 8 bytes.■ MAC type is <code>DL_FDDI</code>.■ SAP length value is -2, meaning the physical address component is followed immediately by a 2-byte SAP component within the DLSAP address.■ Service mode is <code>DL_CLDLS</code>.■ Optional quality of service (QOS) support is not included; as a result, the QOS field values are 0.■ Provider style is <code>DL_STYLE2</code>.■ Version is <code>DL_VERSION_2</code>.

- OPTIONS** Options are not required for normal operation. In special cases, FDDI Station Management (SMT) parameters can be modified by using the `/usr/bin/smtpara` utility (see the driver README files). The `smtpara` utility should be used only by those very familiar with FDDI.
- DIAGNOSTICS** If multiple NICs are installed in the system, the following message may appear on the console and in the `/var/adm/messages` log file:
- ```
skfp: DMA memory allocation failed !
```
- You can avoid this message by tuning the `lomempages` kernel parameter. By default, the value of this parameter is 36 pages. Each SK-FDDI PCI adapter requires nine pages, so increase the value of the `lomempages` parameter in increments of nine pages until all NICs in the system run correctly.
- To modify the value of this parameter to 45 pages, you can, for example, append the `set lomempages=45` line to the `/etc/system` file and reboot the system.
- FILES**
- |                                      |                                                                           |
|--------------------------------------|---------------------------------------------------------------------------|
| <code>/dev/skfp</code>               | Character special device                                                  |
| <code>/dev/skfp<math>x</math></code> | Per-adapter character special device, where $x$ is the adapter <i>ppa</i> |
| <code>/kernel/drv/skfp</code>        | ELF kernel module                                                         |
| <code>/kernel/drv/skfp.conf</code>   | Driver configuration file                                                 |
| <code>/usr/bin/smtpara</code>        | SMT parameter utility                                                     |
| <code>/etc/fddi.cfg</code>           | <code>smtpara</code> configuration file                                   |
| <code>&lt;sys/dlpi.h&gt;</code>      | DLPI definitions                                                          |
- ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Architecture   | IA, SPARC       |

**SEE ALSO** `ifconfig(1M)`, `netstat(1M)`, `attributes(5)`, `dlpi(7P)`

`skfp.txt` README file — Included in the driver package or available from [www.syskonnect.com](http://www.syskonnect.com).

slp(7P)

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                     |                                                                                                                                                                                                                                                                                          |                  |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | slp – Service Location Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                     |                                                                                                                                                                                                                                                                                          |                  |                                                                                                                                                                                                                                                                                                                                                       |
| <b>DESCRIPTION</b>  | <p>The Service Location Protocol (SLP) is a dynamic service discovery protocol that runs on top of the Internet Protocol (IP). The protocol is specified by the IETF standard-track documents <i>RFC 2165</i>, <i>RFC 2608</i>, <i>RFC 2609</i>; the API is documented in <i>RFC 2614</i>.</p> <p>There are two components to the SLP technology. The first is a daemon, <code>slpd(1M)</code>, which coordinates SLP operations. The second is a software library, <code>slp_api(3slp)</code>, through which processes access a public API. Both components are configured by means of the SLP configuration file, <code>slp.conf(4)</code>.</p> <p>The SLP API is useful for two types of processes:</p> <table><tr><td>Client Applications</td><td>Services and service information can be requested from the API. Clients do not need to know the location of a required service, only the type of service, and optionally, the service characteristics. SLP will supply the location and other information to the client through the API.</td></tr><tr><td>Server Processes</td><td>Programs that offer network services use the SLP API to advertise their location as well as other service information. The advertisement can optionally include attributes describing the service. Advertisements are accompanied by a lifetime; when the lifetime expires, the advertisement is flushed, unless it is refreshed prior to expiration.</td></tr></table> <p>API libraries are available for both the C and Java languages.</p> <p>SLP provides the following additional features:</p> <ul style="list-style-type: none"><li>■ <code>slpd(1M)</code> can be configured to function as a transparent directory agent. This feature makes SLP scalable to the enterprise. System administrators can configure directory agents to achieve a number of different strategies for scalability.</li><li>■ SLP service advertising and discovery is performed in scopes. Unless otherwise configured, all discovery and all advertisements are in the scope <i>default</i>. In the case of a larger network, scopes can be used to group services and client systems so that users will only find those services which are physically near them, belong to their department, or satisfy the specified criteria. Administrators can configure these scopes to achieve different service provider strategies.</li><li>■ Services may be registered by proxy through a serialized registration file. This is an alternative to registering services through the API. See <code>slpd.reg(4)</code> for more information.</li></ul> | Client Applications | Services and service information can be requested from the API. Clients do not need to know the location of a required service, only the type of service, and optionally, the service characteristics. SLP will supply the location and other information to the client through the API. | Server Processes | Programs that offer network services use the SLP API to advertise their location as well as other service information. The advertisement can optionally include attributes describing the service. Advertisements are accompanied by a lifetime; when the lifetime expires, the advertisement is flushed, unless it is refreshed prior to expiration. |
| Client Applications | Services and service information can be requested from the API. Clients do not need to know the location of a required service, only the type of service, and optionally, the service characteristics. SLP will supply the location and other information to the client through the API.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                     |                                                                                                                                                                                                                                                                                          |                  |                                                                                                                                                                                                                                                                                                                                                       |
| Server Processes    | Programs that offer network services use the SLP API to advertise their location as well as other service information. The advertisement can optionally include attributes describing the service. Advertisements are accompanied by a lifetime; when the lifetime expires, the advertisement is flushed, unless it is refreshed prior to expiration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                     |                                                                                                                                                                                                                                                                                          |                  |                                                                                                                                                                                                                                                                                                                                                       |



**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Availability        | SUNWslpu        |
| CSI                 | CSI-enabled     |
| Interface Stability | Standard        |
| MT-Level            | MT-Safe         |

**SEE ALSO** `slpd(1m)`, `slp_api(3slp)`, `slp.conf(4)`, `slpd.reg(4)`, `attributes(5)`

Guttman, E., Perkins, C., Veizades, J., and Day, M., *RFC 2608, Service Location Protocol, Version 2*, The Internet Society, June 1999.

Guttman, E., Perkins, C., and Kempf, J., *RFC 2609, Service Templates and Service Schemes*, The Internet Society, June 1999.

Kempf, J. and Guttman, E., *RFC 2614, An API for Service Location*, The Internet Society, June 1999.

Veizades, J., Guttman, E., Perkins, C., and Kaplan, S., *RFC 2165, Service Location Protocol*, Network Working Group, 1997.

smbus(7D)

**NAME** | smbus – System Management Bus controller driver

**DESCRIPTION** | The smbus driver is a I2C (Inter IC) nexus driver that allows the system to communicate with various system component chips. SMBus is a two-wire control bus based on the I2C protocol through which systems can communicate with various I2C devices connected to the bus.

| The smbus driver supports byte and block level transfer based on interrupt and polled mode.

**FILES** | /platform/sun4u/kernel/drv/sparcv9/smbus  
64 bit ELF kernel module

**ATTRIBUTES** | See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability   | SUNWcarx        |
| Architecture   | SPARC           |

**SEE ALSO** | `attributes(5)`

*Writing Device Drivers*

*System Management Bus (SMBus) Specification 2.0* — SBS Implementation Forum

*The I2C Bus and How To Use It* —Philips Semiconductor Document # 98-8080-575-01

socal(7D)

**NAME** socal – Serial Optical Controller for Fibre Channel Arbitrated Loop (SOC+) device driver

**SYNOPSIS** socal@sbus-slot, 0

**DESCRIPTION** The Fibre Channel Host Bus Adapter is an SBus card which implements two full duplex Fibre Channel interfaces. Each Fibre Channel interface can connect to a Fibre Channel Arbitrated Loop (FC-AL).

The socal device driver is a nexus driver and implements portions of the FC-2 and FC-4 layers of FC-AL.

**FILES** /kernel/drv/socal ELF Kernel Module

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Architecture   | SPARC           |

**SEE ALSO** sbus(4), sf(7D), ssd(7D)

*Writing Device Drivers*

*ANSI X3.230-1994, Fibre Channel Physical and Signalling Interface (FC-PH)*

*ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL)*

*Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)*

**DIAGNOSTICS** The messages described below may appear on system console in addition to being logged.

On the console, these messages are preceded by:

socal%d: port %a where %d is the instance number of the socal controller and %a is the port on the host adapter.

Fibre Channel Loop is ONLINE

The Fibre Channel loop is now online.

Fibre Channel Loop is OFFLINE

The Fibre Channel loop is now offline.

attach failed: device in slave-only slot.

Move soc+ card to another slot.

attach failed: bad soft state.

Driver did not attach, devices will be inaccessible.

attach failed: unable to alloc xport struct.

Driver did not attach, devices will be inaccessible.

socal(7D)

attach failed: unable to map eeprom

Driver was unable to map device memory; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to map XRAM

Driver was unable to map device memory; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to map registers

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to access status register

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to install interrupt handler

Driver was not able to add the interrupt routine to the kernel. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to access host adapter XRAM

Driver was unable to access device RAM; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to write host adapter XRAM

Driver was unable to write device RAM; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: read/write mismatch in XRAM

Driver was unable to verify device RAM; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | sockio – ioctls that operate directly on sockets                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>SYNOPSIS</b>    | <code>#include &lt;sys/sockio.h&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>DESCRIPTION</b> | <p>The <code>ioctls</code> listed in this manual page apply directly to sockets, independent of any underlying protocol. The <code>setsockopt()</code> call (see <code>getsockopt(3SOCKET)</code>) is the primary method for operating on sockets, rather than on the underlying protocol or network interface. <code>ioctls</code> for a specific network interface or protocol are documented in the manual page for that interface or protocol.</p> <p><code>SIOCSPGRP</code>      The argument is a pointer to an <code>int</code>. Set the process-group ID that will subsequently receive <code>SIGIO</code> or <code>SIGURG</code> signals for the socket referred to by the descriptor passed to <code>ioctl</code> to the value of that <code>int</code>. The argument must be either positive (in which case it must be a process ID) or negative (in which case it must be a process group).</p> <p><code>SIOCGPGRP</code>      The argument is a pointer to an <code>int</code>. Set the value of that <code>int</code> to the process-group ID that is receiving <code>SIGIO</code> or <code>SIGURG</code> signals for the socket referred to by the descriptor passed to <code>ioctl</code>.</p> <p><code>SIOCCATMARK</code>    The argument is a pointer to an <code>int</code>. Set the value of that <code>int</code> to 1 if the read pointer for the socket referred to by the descriptor passed to <code>ioctl</code> points to a mark in the data stream for an out-of-band message. Set the value of that <code>int</code> to 0 if the read pointer for the socket referred to by the descriptor passed to <code>ioctl</code> does not point to a mark in the data stream for an out-of-band message.</p> |
| <b>SEE ALSO</b>    | <code>ioctl(2)</code> , <code>getsockopt(3SOCKET)</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## sppptun(7M)

|                    |                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | sppptun – PPP tunneling pseudo-driver                                                                                                                        |
| <b>SYNOPSIS</b>    | /dev/sppptun                                                                                                                                                 |
| <b>DESCRIPTION</b> | The /dev/sppptun pseudo-driver provides an interface for tunneling PPP sessions. This interface provides PPP over Ethernet (PPPoE) service with Solaris PPP. |
| <b>FILES</b>       | /dev/sppptun Solaris PPP tunneling device driver.                                                                                                            |
| <b>SEE ALSO</b>    | pppoe(1M), pppoe(1M), sppptun(1M)<br><i>RFC 2516 — A Method for Transmitting PPP Over Ethernet (PPPoE)</i> . Mamakos, et. al. February 1999.                 |

| <b>NAME</b>                              | spwr – SMC EtherPower II 10/100 (9432) Ethernet device driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                           |                |                 |              |     |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|----------------|-----------------|--------------|-----|
| <b>SYNOPSIS</b>                          | <code>/dev/spwr</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                           |                |                 |              |     |
| <b>DESCRIPTION</b>                       | <p>The <code>spwr</code> Ethernet driver is a multi-threaded, loadable, clonable, STREAMS hardware driver supporting the connectionless Data Link Provider Interface, <code>dlpi(7P)</code>, over SMC EtherPower II 10/100 controllers. Multiple EtherPower II controllers installed within the system are supported by the driver. The <code>spwr</code> driver provides basic support for the SMC EtherPower II hardware. Functions include chip initialization, frame transmit and receive, multicast support, and error recovery and reporting.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                           |                |                 |              |     |
| <b>APPLICATION PROGRAMMING INTERFACE</b> | <p>The cloning character-special device <code>/dev/spwr</code> is used to access all SMC EtherPower II devices installed within the system.</p> <p>The <code>spwr</code> driver is dependent on <code>/kernel/misc/gld</code>, a loadable kernel module that provides the <code>spwr</code> driver with the DLPI and STREAMS functionality required of a LAN driver. See <code>gld(7d)</code> for more details on the primitives supported by the driver.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> from the user are as follows:</p> <ul style="list-style-type: none"> <li>■ The maximum SDU is 1500 ( <code>ETHERMTU</code>).</li> <li>■ The minimum SDU is 0. The <code>spwr</code> driver will pad to the mandatory 60-octet minimum packet size.</li> <li>■ The DLSAP address length is 8.</li> <li>■ The MAC type is <code>DL_ETHER</code>.</li> <li>■ The SAP length value is -2, meaning the physical address component is followed immediately by a 2-byte SAP component within the DLSAP address.</li> <li>■ The broadcast address value is Ethernet/IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>).</li> </ul> |                           |                |                 |              |     |
| <b>FILES</b>                             | <code>/dev/spwr</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Character special device. |                |                 |              |     |
|                                          | <code>/kernel/drv/spwr.conf</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Driver configuration file |                |                 |              |     |
| <b>ATTRIBUTES</b>                        | See <code>attributes(5)</code> for descriptions of the following attributes:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                           |                |                 |              |     |
|                                          | <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>x86</td> </tr> </tbody> </table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                           | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Architecture | x86 |
| ATTRIBUTE TYPE                           | ATTRIBUTE VALUE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                           |                |                 |              |     |
| Architecture                             | x86                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                           |                |                 |              |     |
| <b>SEE ALSO</b>                          | <code>attributes(5)</code> , <code>dlpi(7P)</code> , <code>gld(7D)</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                           |                |                 |              |     |

ssd(7D)

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|--------------------|--------------------|---------------------------------------------------------|---------------------|---------------------------------|---------------------|-------------------|------------------|------------------------|---------------------|-----------------------------------------------------------|
| <b>NAME</b>                      | ssd – Fibre Channel Arbitrated Loop disk device driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <b>SYNOPSIS</b>                  | <i>ssd@port , target : partition</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <b>DESCRIPTION</b>               | <p>The <code>ssd</code> driver supports Fibre Channel disk devices.</p> <p>The specific type of each disk is determined by the SCSI inquiry command and reading the volume label stored on block 0 of the drive. The volume label describes the disk geometry and partitioning; it must be present or the disk cannot be mounted by the system.</p> <p>The block-files access the disk using the system's normal buffering mechanism and are read and written without regard to physical disk records. A "raw" interface provides for direct transmission between the disk and the read or write buffer. A single read or write call usually results in one I/O operation; raw I/O is therefore more efficient when many bytes are transmitted. Block file names are found in <code>/dev/dsk</code>; the names of the raw files are found in <code>/dev/rdisk</code>.</p> <p>I/O requests (such as <code>lseek(2)</code>) to the SCSI disk must have an offset that is a multiple of 512 bytes (<code>DEV_BSIZE</code>), or the driver returns an <code>EINVAL</code> error. If the transfer length is not a multiple of 512 bytes, the transfer count is rounded up by the driver.</p> <p>Partition 0 is normally used for the root file system on a disk, with partition 1 as a paging area (for example, <code>swap</code>). Partition 2 is used to back up the entire disk. Partition 2 normally maps the entire disk and may also be used as the mount point for secondary disks in the system. The rest of the disk is normally partition 6. For the primary disk, the user file system is located here.</p> <p>The device has associated error statistics. These must include counters for hard errors, soft errors and transport errors. Other data may be implemented as required.</p> |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <b>DEVICE STATISTICS SUPPORT</b> | <p>The device maintains I/O statistics for the device and for partitions allocated for that device. For each device/partition, the driver accumulates reads, writes, bytes read, and bytes written. The driver also initiates hi-resolution time stamps at queue entry and exit points to enable monitoring of residence time and cumulative residence-length product for each queue.</p> <p>Not all device drivers make per-partition IO statistics available for reporting. <code>ssd</code> and <code>sd(7D)</code> per-partition statistics are enabled by default but may be disabled in their configuration files.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <b>IOCTLS</b>                    | Refer to <code>dkio(7I)</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <b>ERRORS</b>                    | <table><tr><td><code>EACCES</code></td><td>Permission denied.</td></tr><tr><td><code>EBUSY</code></td><td>The partition was opened exclusively by another thread.</td></tr><tr><td><code>EFAULT</code></td><td>The argument was a bad address.</td></tr><tr><td><code>EINVAL</code></td><td>Invalid argument.</td></tr><tr><td><code>EIO</code></td><td>An I/O error occurred.</td></tr><tr><td><code>ENOTTY</code></td><td>The device does not support the requested ioctl function.</td></tr></table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <code>EACCES</code> | Permission denied. | <code>EBUSY</code> | The partition was opened exclusively by another thread. | <code>EFAULT</code> | The argument was a bad address. | <code>EINVAL</code> | Invalid argument. | <code>EIO</code> | An I/O error occurred. | <code>ENOTTY</code> | The device does not support the requested ioctl function. |
| <code>EACCES</code>              | Permission denied.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <code>EBUSY</code>               | The partition was opened exclusively by another thread.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <code>EFAULT</code>              | The argument was a bad address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <code>EINVAL</code>              | Invalid argument.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <code>EIO</code>                 | An I/O error occurred.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |
| <code>ENOTTY</code>              | The device does not support the requested ioctl function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                     |                    |                    |                                                         |                     |                                 |                     |                   |                  |                        |                     |                                                           |



|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|---------------------------|--------------------------------|-------------|----------------------------------|-----------|-----------------|-----------------------------------------|-----------------|----------------------------------------------------|-----------------|-------------------|-----------------|--------------------------|
|                                  | <p>ENXIO            When returned during <code>open(2)</code>, this error indicates the device does not exist.</p> <p>EROFS            The device is a read-only device.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <b>CONFIGURATION</b>             | <p>You configure the <code>ssd</code> driver by defining properties in the <code>ssd.conf</code> file. The <code>ssd</code> driver supports the following properties:</p> <p><code>enable-partition-kstats</code><br/>         The default value is 1, which causes partition IO statistics to be maintained. Set this value to zero to prevent the driver from recording partition statistics. This slightly reduces the CPU overhead for IO, minimizes the amount of <code>sar(1)</code> data collected and makes these statistics unavailable for reporting by <code>iostat(1M)</code> even though the <code>-p/-P</code> option is specified. Regardless of this setting, disk IO statistics are always maintained.</p>                                                                                                                                                                                                                                         |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <b>FILES</b>                     | <table border="0"> <tr> <td><code>ssd.conf</code></td> <td>Driver configuration file</td> </tr> <tr> <td><code>/dev/dsk/cntndnsn</code></td> <td>block files</td> </tr> <tr> <td><code>/dev/rdisk/cntndnsn</code></td> <td>raw files</td> </tr> <tr> <td><code>cn</code></td> <td>is the controller number on the system.</td> </tr> <tr> <td><code>tn</code></td> <td>7-bit disk loop identifier, such as switch setting</td> </tr> <tr> <td><code>dn</code></td> <td>SCSI lun <i>n</i></td> </tr> <tr> <td><code>sn</code></td> <td>partition <i>n</i> (0-7)</td> </tr> </table>                                                                                                                                                                                                                                                                                                                                                                                  | <code>ssd.conf</code> | Driver configuration file | <code>/dev/dsk/cntndnsn</code> | block files | <code>/dev/rdisk/cntndnsn</code> | raw files | <code>cn</code> | is the controller number on the system. | <code>tn</code> | 7-bit disk loop identifier, such as switch setting | <code>dn</code> | SCSI lun <i>n</i> | <code>sn</code> | partition <i>n</i> (0-7) |
| <code>ssd.conf</code>            | Driver configuration file                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <code>/dev/dsk/cntndnsn</code>   | block files                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <code>/dev/rdisk/cntndnsn</code> | raw files                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <code>cn</code>                  | is the controller number on the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <code>tn</code>                  | 7-bit disk loop identifier, such as switch setting                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <code>dn</code>                  | SCSI lun <i>n</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <code>sn</code>                  | partition <i>n</i> (0-7)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <b>SEE ALSO</b>                  | <p><code>sar(1)</code>, <code>format(1M)</code>, <code>iostat(1M)</code>, <code>ioctl(2)</code>, <code>lseek(2)</code>, <code>open(2)</code>, <code>read(2)</code>, <code>write(2)</code>, <code>driver.conf(4)</code>, <code>cdio(7I)</code>, <code>dkio(7I)</code></p> <p><i>ANSI Small Computer System Interface-2 (SCSI-2)</i></p> <p><i>ANSI X3.272-1996, Fibre Channel Arbitrated Loop (FC-AL)</i></p> <p><i>Fibre Channel - Private Loop SCSI Direct Attach (FC-PLDA)</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |
| <b>DIAGNOSTICS</b>               | <p>Error for command '<code>&lt;command name&gt;</code>' Error Level: Fatal Requested Block <code>&lt;n&gt;</code>, Error Block: <code>&lt;m&gt;</code>, Sense Key: <code>&lt;sense key name&gt;</code>, Vendor '<code>&lt;vendor name&gt;</code>':<br/>         ASC = <code>0x&lt;a&gt;</code> (<code>&lt;ASC name&gt;</code>), ASCQ = <code>0x&lt;b&gt;</code>, FRU = <code>0x&lt;c&gt;</code></p> <p>The command indicated by <code>&lt;command name&gt;</code> failed. The Requested Block is the block where the transfer started and the Error Block is the block that caused the error. Sense Key, ASC, and ASCQ information is returned by the target in response to a request sense command.</p> <p>Check Condition on REQUEST SENSE</p> <p>A REQUEST SENSE command completed with a check condition. The original command will be retried a number of times.</p> <p>Label says <code>&lt;m&gt;</code> blocks Drive says <code>&lt;n&gt;</code> blocks</p> |                       |                           |                                |             |                                  |           |                 |                                         |                 |                                                    |                 |                   |                 |                          |

ssd(7D)

There is a discrepancy between the label and what the drive returned on the READ CAPACITY command.

Not enough sense information

The request sense data was less than expected.

Request Sense couldn't get sense data

The REQUEST SENSE command did not transfer any data.

Reservation Conflict

The drive was reserved by another initiator.

SCSI transport failed: reason 'xxxx' : {retrying|giving up}

The host adapter has failed to transport a command to the target for the reason stated. The driver will either retry the command or, ultimately, give up.

Unhandled Sense Key <#>

The REQUEST SENSE data included an invalid sense key.

Unit not Ready. Additional sense code 0x<#>

The drive is not ready.

corrupt label - bad geometry

The disk label is corrupted.

corrupt label - label checksum failed

The disk label is corrupted.

corrupt label - wrong magic number

The disk label is corrupted.

device busy too long

The drive returned busy during a number of retries.

disk not responding to selection

The drive was probably powered down or died.

i/o to invalid geometry

The geometry of the drive could not be established.

incomplete read/write - retrying/giving up

There was a residue after the command completed normally.

logical unit not ready

The drive is not ready.

no bp for disk label

A bp with consistent memory could not be allocated.

no mem for property

Free memory pool exhausted.

no memory for disk label

Free memory pool exhausted.

no resources for dumping

A packet could not be allocated during dumping.

offline

Drive went offline; probably powered down.

requeue of command fails<n>

Driver attempted to retry a command and experienced a transport error.

ssdrestart transport failed <n>

Driver attempted to retry a command and experienced a transport error.

transfer length not modulo <n>

Illegal request size.

transport rejected <n>

Host adapter driver was unable to accept a command.

unable to read label

Failure to read disk label.

unit does not respond to selection

Drive went offline; probably powered down.

## st(7D)

|                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                                              | st – driver for SCSI tape devices                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>SYNOPSIS</b>                                          | <code>st@target, lun: [l, m, h, c, u] [b] [n]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>DESCRIPTION</b>                                       | <p>The <code>st</code> device driver provides a standard interface to various SCSI tape devices. See <a href="#">mtio(7I)</a> for details.</p> <p>To determine if the <code>st</code> device driver supports your tape device, SPARC users should enter the following on a command line:</p> <pre>% strings /kernel/drv/sparcv9   grep -i &lt;tape device name&gt; (e.g. "Quantam").</pre> <p>x86 users can do the following to determine if the <code>st</code> device driver supports a particular tape device:</p> <pre>% strings /kernel/drv/st   grep -i &lt;tape device name&gt; (e.g. "Emulex").</pre> <p>The driver can be opened with either <code>rewind on close</code> or <code>no rewind on close</code> options. It can also be opened with the <code>O_NDELAY</code> (see <a href="#">open(2)</a>) option when there is no tape inserted in the drive. A maximum of four tape formats per device are supported (see <a href="#">FILES</a> below). The tape format is specified using the device name. (Tape format is also referred to as tape density).</p> <p>The driver reserves the tape drive upon open and releases it at close for use in multi-initiator environments. Refer to the <code>MTIOCRESERVE</code> and <code>MTIOCRELEASE</code> ioctls in <a href="#">mtio(7I)</a> for information about how to allow a tape drive to remain reserved upon close. See the flag options below for information about disabling this feature.</p> <p>If the tape drive is opened in <code>O_NDELAY</code> mode, no reservation will occur during the open, as per the POSIX standard (see <a href="#">standards(5)</a>). However, before the first tape operation or I/O occurs, a reservation will occur to provide reserve/release functionality.</p> |
| <b>Persistent Errors and Asynchronous Tape Operation</b> | The <code>st</code> driver now supports persistent errors (see <a href="#">mtio(7I)</a> ) and asynchronous tape operations (see <a href="#">mtio(7I)</a> , <a href="#">aioread(3AIO)</a> , and <a href="#">aiowrite(3AIO)</a> ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Read Operation</b>                                    | <p>If the driver is opened for reading in a different format than the tape is written in, the driver overrides the user-selected format. For example, if a 1/4" cartridge tape is written in QIC-24 format and opened for reading in QIC-150, the driver detects a read failure on the first read and automatically switches to QIC-24 to read the data.</p> <p>Note that if the low density format is used, no indication is given that the driver has overridden the format you selected. Other formats issue a warning message to inform you of an overridden format selection. Some devices automatically perform this function and do not require driver support (1/2" reel tape drive, for example).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Write Operation</b>                                   | Writing from the beginning of tape is performed in the user-specified format. The original tape format is used for appending onto previously written tapes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Tape Configuration**

The `st` tape driver has a built-in configuration table for all Sun supported tape drives. To support the addition of third party tape devices or to override a built-in configuration, device information can be supplied in `st.conf` as global properties that apply to each node, or as properties that are applicable to one node only. The `st` driver looks for the property called "tape-config-list." The value of this property is a list of triplets, where each triplet consists of three strings.

The formal syntax is:

```
tape-config-list = <triplet> [, <triplet> *];
where
<triplet> := <vid+pid>, <pretty print>, <data-property-name>
and
<data-property-name> = <version>, <type>, <bsize>,
 <options>, <number of densities>,
 <density> [, <density>*], <default-density>;
or
<data-property-name> = <version 2>, <type>, <bsize>,
 <options>, <number of densities>,
 <density> [, <density>*], <default-density>,
 <non-motion time-out>, <I/O time-out>,
 <rewind time-out>, <space time-out>,
 <load time-out>, <unload time-out>,
 <erase time-out>;
```

A semicolon (;) is used to terminate a prototype devinfo node specification. Individual elements listed within the specification should not be separated by a semicolon. (Refer to `driver.conf(4)` for more information.)

<vid+pid> is the string that is returned by the tape device on a SCSI inquiry command. This string may contain any character in the range 0x20-0x7e. Characters such as "" (double quote) or "'" (single quote), which are not permitted in property value strings, are represented by their octal equivalent (for example, \042 and \047). Trailing spaces may be truncated.

<pretty print> is used to report the device on the console. This string may have zero length, in which case the <vid+pid> will be used to report the device.

<data-property-name> is the name of the property which contains all the tape configuration values (such as <type>, <bsize>, etc.) corresponding for the tape drive for the specified <vid+pid>.

<version> is a version number and should be 1 or 2. In the future, higher version numbers may be used to allow for changes in the syntax of the <data-property-name> value list.

<type> is a type field. Valid types are defined in `/usr/include/sys/mtio.h`. For third party tape configuration, the following generic types are recommended:

st(7D)

---

|              |      |
|--------------|------|
| MT_ISQIC     | 0x32 |
| MT_ISREEL    | 0x33 |
| MT_ISDAT     | 0x34 |
| MT_IS8MM     | 0x35 |
| MT_ISOTHER   | 0x36 |
| MT_ISTAND25G | 0x37 |
| MT_ISDLT     | 0x38 |
| MT_ISSTK9840 | 0x39 |
| MT_ISBMDLT1  | 0x3A |
| MT_LTO       | 0x3B |

---

<bsize> is the preferred block size of the tape device. The value should be 0 for variable block size devices.

<options> is a bit pattern representing the devices, as defined in /usr/include/sys/scsi/targets/stdef.h. Valid flags for tape configuration are:

---

|                         |         |
|-------------------------|---------|
| ST_VARIABLE             | 0x0001  |
| ST_QIC                  | 0x0002  |
| ST_REEL                 | 0x0004  |
| ST_BSF                  | 0x0008  |
| ST_BSR                  | 0x0010  |
| ST_LONG_ERASE           | 0x0020  |
| ST_AUTODEN_OVERRIDE     | 0x0040  |
| ST_NOBUF                | 0x0080  |
| ST_KNOWS_EOD            | 0x0200  |
| ST_UNLOADABLE           | 0x0400  |
| ST_SOFT_ERROR_REPORTING | 0x0800  |
| ST_LONG_TIMEOUTS        | 0x1000  |
| ST_NO_RECSIZE_LIMIT     | 0x8000  |
| ST_MODE_SEL_COMP        | 0x10000 |

---

|                                      |          |
|--------------------------------------|----------|
| ST_NO_RESERVE_RELEASE                | 0x20000  |
| ST_READ_IGNORE_ILI                   | 0x40000  |
| ST_READ_IGNORE_EOFs                  | 0x80000  |
| ST_SHORT_FILEMARKS                   | 0x100000 |
| ST_EJECT_TAPE_ON_CHANGER_FAILURE     | 0x200000 |
| ST_RETRY_ON_RECOVERED_DEFERRED_ERROR | 0x400000 |

**ST\_VARIABLE**

The flag indicates the tape device supports variable length record sizes.

**ST\_QIC**

The flag indicates a Quarter Inch Cartridge (QIC) tape device.

**ST\_REEL**

The flag indicates a 1/2-inch reel tape device.

**ST\_BSF**

If flag is set, the device supports backspace over EOF marks (bsf - see mt(1)).

**ST\_BSR**

If flag is set, the tape device supports the backspace record operation (bsr - see mt(1)). If the device does not support bsr, the st driver emulates the action by rewinding the tape and using the forward space record (fsf) operation to forward the tape to the correct file. The driver then uses forward space record (fsr - see mt(1)) to forward the tape to the correct record.

**ST\_LONG\_ERASE**

The flag indicates the tape device needs a longer time than normal to erase.

**ST\_AUTODEN\_OVERRIDE**

The auto-density override flag. The device is capable of determining the tape density automatically without issuing a "mode-select"/"mode-sense command."

**ST\_NOBUF**

The flag disables the device's ability to perform buffered writes. A buffered write occurs when the device acknowledges the completion of a write request after the data has been written to the device's buffer, but before all of the data has been written to the tape.

**ST\_KNows\_EOD**

If flag is set, the device can determine when EOD (End of Data) has been reached. When this flag is set, the st driver uses fast file skipping. Otherwise, file skipping happens one file at a time.

**ST\_UNLOADABLE**

The flag indicates the device will not complain if the st driver is unloaded and loaded again (see modload(1M) and modunload(1M)). That is, the driver will return the correct inquiry string.

st(7D)

**ST\_SOFT\_ERROR\_REPORTING**

The flag indicates the tape device will perform a “request sense” or “log sense” command when the device is closed. Currently, only Exabyte and DAT drives support this feature.

**ST\_LONG\_TIMEOUTS**

The flag indicates the tape device requires timeouts that are five times longer than usual for normal operation.

**ST\_NO\_RECSIZE\_LIMIT**

The flag applies to variable-length tape devices. If this flag is set, the record size is not limited to a 64 Kbyte record size. The record size is only limited by the smaller of either the record size supported by the device or the maximum DMA transfer size of the system. (Refer to `Large Record Sizes` and `WARNINGS`.) The maximum block size that will not be broken into smaller blocks can be determined from the `mt_bf` returned from the `MTIOCGGET` ioctl(). This number is the lesser of the upper block limit returned by the drive from `READ BLOCK LIMITS` command and the `dma-max` property set by the Host Bus Adapter (HBA) to which the drive is attached.

**ST\_MODE\_SEL\_COMP**

If the `ST_MODE_SEL_COMP` flag is set, the driver determines which of the two mode pages the device supports for selecting or deselecting compression. It first tries the Data Compression mode page (0x0F); if this fails, it tries the Device Configuration mode page (0x10). Some devices, however, may need a specific density code for selecting or deselecting compression. Please refer to the device specific SCSI manual. When the flag is set, compression is enabled only if the “c” or “u” device is used. Note that when the lower 2 densities of a drive are identically configured and the upper 2 densities are identically configured, but the lower and upper differ from each other and `ST_MODE_SEL_COMP` is set, the “m” node sets compression *on* for the lower density code (for example, 0x42) and the “c” and “u” nodes set compression *on* for the higher density (for example, 0x43). For any other device densities, compression is disabled.

**ST\_NO\_RESERVE\_RELEASE**

The `ST_NO_RESERVE_RELEASE` flag disables the use of reserve on open and release on close. If an attempt to use a ioctl of `MTRESERVE` or `MTRERELEASE` on a drive with this flag set, it will return an error of `ENOTTY` (inappropriate ioctl for device).

**ST\_READ\_IGNORE\_ILI**

The `ST_READ_IGNORE_ILI` flag is applicable only to variable block devices which support the SILI bit option. The `ST_READ_IGNORE_ILI` flag indicates that SILI (suppress incorrect length indicator) bit will be set during reads. When this flag is set, short reads (requested read size is less than the record size on the tape) will be successful and the number of bytes transferred will be equal to the record size on the tape. The tape will be positioned at the start of the next record skipping over the extra data (the remaining data has been lost). Long reads (requested read size is more than the record size on the tape) will see a large performance gain when this flag is set, due to overhead reduction. When this flag is not set, short reads will return an error of `ENOMEM`.



**ST\_READ\_IGNORE\_EOFs**

The **ST\_READ\_IGNORE\_EOFs** flag is applicable only to 1/2" Reel Tape drives and when performing consecutive reads only. It should not be used for any other tape command. Usually End-of-recorded-media (EOM) is indicated by two EOF marks on 1/2" tape and application cannot read past EOM. When this flag is set, two EOF marks no longer indicate EOM allowing applications to read past two EOF marks. In this case it is the responsibility of the application to detect end-of-recorded-media (EOM). When this flag is set, tape operations (like MTEOM) which positions the tape at end-of-recorded-media will fail since detection of end-of-recorded-media (EOM) is to be handled by the application. This flag should be used when backup applications have embedded double filemarks between files.

**ST\_SHORT\_FILEMARKS**

The **ST\_SHORT\_FILEMARKS** flag is applicable only to EXABYTE 8mm tape drives which supports short filemarks. When this flag is set, short filemarks will be used for writing filemarks. Short filemarks could lead to tape incompatible with some otherwise compatible device. By default long filemarks will be used for writing filemarks.

**ST\_EJECT\_TAPE\_ON\_CHANGER\_FAILURE**

If **ST\_EJECT\_TAPE\_ON\_CHANGER\_FAILURE** flag is set, the tape will be ejected automatically if the tape cartridge is trapped in the medium due to positioning problems of the medium changer.

The following ASC/ASCQ keys are defined to the reasons for causing tape ejection if **ST\_EJECT\_TAPE\_ON\_CHANGER\_FAILURE** option is set to 0x200000:

Sense ASC/ASCQ Description

Key

4 15/01 Mechanical Failure

4 44/00 Internal Target Failure

2 53/00 Media Load or Eject Failed

4 53/00 Media Load or Eject Failed

4 53/01 Unload Tape Failure

**ST\_RETRY\_ON\_RECOVERED\_DEFERRED\_ERROR**

If **ST\_RETRY\_ON\_RECOVERED\_DEFERRED\_ERROR** flag is set, the st driver will retry the last write if this cmd caused a check condition with error code 0x71 and sense code 0x01. Some tape drives, notably the IBM 3090, require this option.

<number of densities> is the number of densities specified. Each tape drive can support up to four densities. The value entered should therefore be between 1 and 4; if less than 4, the remaining densities will be assigned a value of 0x0.

st(7D)

<density> is a single-byte hexadecimal number. It can either be found in the device specification manual or be obtained from the device vendor.

<default-density> has a value between 0 and (<number of densities> - 1).

<non-motion time-out> Time in seconds that the drive should be able to perform any SCSI command that doesn't require tape to be moved. This includes mode sense, mode select, reserve, release, read block limits, and test unit ready.

<I/O time-out> Time in seconds to perform data transfer I/O to or from tape including worst case error recovery.

<rewind time-out> Time in seconds to rewind from anywhere on tape to BOT including worst case recovery forcing buffered write data to tape.

<space time-out> Time in seconds to space to any file, block or end of data on tape. Including worst case when any form of cataloging is invalid.

<load time-out> Time in seconds to load tape and be ready to transfer first block. This should include worst case recovery reading tape catalog or drive specific operations done at load.

<unload time-out> Time in seconds to unload tape. Should include worst case time to write to catalog, unthread, and tape cartridge unloading. Also should include worst case time for any drive specific operations that are preformed at unload. Should not include rewind time as the driver rewinds tape before issuing the unload.

<erase time-out> Time in seconds to preform a full (BOT to EOT) erase of longest medium with worst case error recovery.

#### Device Statistics Support

Each device maintains I/O statistics both for the device and for each partition allocated on that device. For each device/partition, the driver accumulates reads, writes, bytes read, and bytes written. The driver also takes hi-resolution time stamps at queue entry and exit points, which facilitates monitoring the residence time and cumulative residence-length product for each queue.

Each device also has error statistics associated with it. These must include counters for hard errors, soft errors and transport errors. Other data may be implemented as required.

#### IOCTLS

The behavior of SCSI tape positioning ioctls is the same across all devices which support them. (Refer to `mtio(7I)`.) However, not all devices support all ioctls. The driver returns an `ENOTTY` error on unsupported ioctls.

The retension ioctl only applies to 1/4" cartridge tape devices. It is used to restore tape tension, thus improving the tape's soft error rate after extensive start-stop operations or long-term storage.

In order to increase performance of variable-length tape devices (particularly when they are used to read/write small record sizes), two operations in the `MTIOCTOP` ioctl, `MTSRSZ` and `MTGRSZ`, can be used to set and get fixed record lengths. The ioctl also

works with fixed-length tape drives which allow multiple record sizes. The min/max limits of record size allowed on a driver are found by using a SCSI-2 READ BLOCK LIMITS command to the device. If this command fails, the default min/max record sizes allowed are 1 byte and 63k bytes. An application that needs to use a different record size opens the device, sets the size with the MTSRSZ ioctl, and then continues with I/O. The scope of the change in record size remains until the device is closed. The next open to the device resets the record size to the default record size (retrieved from `st.conf`).

Note that the error status is reset by the MTIOCGET get status ioctl call or by the next read, write, or other ioctl operation. If no error has occurred (sense key is 0), the current file and record position is returned.

|               |        |                                                                                                                                                                                                           |
|---------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ERRORS</b> | EACCES | The driver is opened for write access and the tape is write-protected or the tape unit is reserved by another host.                                                                                       |
|               | EBUSY  | The tape drive is in use by another process. Only one process can use the tape drive at a time. The driver will allow a grace period for the other process to finish before reporting this error.         |
|               | EINVAL | The number of bytes read or written is not a multiple of the physical record size (fixed-length tape devices only).                                                                                       |
|               | EIO    | During opening, the tape device is not ready because either no tape is in the drive, or the drive is not on-line. Once open, this error is returned if the requested I/O transfer could not be completed. |
|               | ENOTTY | This indicates that the tape device does not support the requested ioctl function.                                                                                                                        |
|               | ENXIO  | During opening, the tape device does not exist.                                                                                                                                                           |
|               | ENOMEM | This indicates that the record size on the tape drive is more than the requested size during read operation.                                                                                              |

#### **EXAMPLES** **EXAMPLE 1** Global tape-config list property

The following is an example of a global `tape-config-list` property:

```
tape-config-list =
"Magic DAT", "Magic 4mm Helical Scan", "magic-data",
"Major Appliance", "Major Appliance Tape", "major-tape";

magic-data = 1,0x34,1024,0x1639,4,0,0x8c,0x8c,0x8c,3;
major-tape = 2,0x3c,0,0x18619,4,0x0,0x0,0x0,0x0,
 3,0,0,30,120,0,0,36000;

name="st" class="scsi"
target=0 lun=0;
```

st(7D)

**EXAMPLE 1** Global tape-config list property (Continued)

```
name="st" class="scsi"
 target=1 lun=0;
name="st" class="scsi"
 target=2 lun=0;
.
.
name="st" class="scsi"
 target=6 lun=0;
```

**EXAMPLE 2** Tape-config-list property applicable to target 2 only

The following is an example of a tape-config-list property applicable to target 2 only:

```
name="st" class="scsi"
 target=0 lun=0;
name="st" class="scsi"
 target=1 lun=0;
name="st" class="scsi"
 target=2 lun=0
 tape-config-list =
 "Magic DAT", "Magic 4mm Helical Scan", "magic-data"
 magic-data = 1,0x34,1024,0x1639,4,0,0x8c,0x8c,0x8c,3;
name="st" class="scsi"
 target=3 lun=0;
.
.
name="st" class="scsi"
 target=6 lun=0;
```

**Large Record Sizes**

To support applications such as seismic programs that require large record sizes, the flag `ST_NO_RECSIZE_LIMIT` must be set in drive option in the configuration entry. A SCSI tape drive that needs to transfer large records should OR this flag with other flags in the 'options' field in `st.conf`. (Refer to Tape Configuration.) By default, this flag is set for the built-in config entries of Archive DAT and Exabyte drives.

If this flag is set, the `st` driver issues a SCSI-2 `READ BLOCK LIMITS` command to the device to determine the maximum record size allowed by it. If the command fails, `st` continues to use the maximum record sizes mentioned in the [mtio\(7I\)](#) man page.

If the command succeeds, `st` restricts the maximum transfer size of a variable-length device to the minimum of that record size and the maximum DMA size that the host adapter can handle. Fixed-length devices are bound by the maximum DMA size allocated by the machine. Note that tapes created with a large record size may not be readable by earlier releases or on other platforms.

(Refer to the `WARNINGS` section for more information.)

**EOT Handling**

The Emulex drives have only a physical end of tape (PEOT); thus it is not possible to write past EOT. All other drives have a logical end of tape (LEOT) before PEOT to guarantee flushing the data onto the tape. The amount of storage between LEOT and PEOT varies from less than 1 Mbyte to about 20 Mbyte, depending on the tape drive.

If EOT is encountered while writing an Emulex, no error is reported but the number of bytes transferred is 0 and no further writing is allowed. On all other drives, the first write that encounters EOT will return a short count or 0. If a short count is returned, then the next write will return 0. After a zero count is returned, the next write returns a full count or short count. A following write returns 0 again. It is important that the number and size of trailer records be kept as small as possible to prevent data loss. Therefore, writing after EOT is not recommended.

Reading past EOT is transparent to the user. Reading is stopped only by reading EOF's. For 1/2" reel devices, it is possible to read off the end of the reel if one reads past the two file marks which mark the end of recorded media.

**FILES**

/kernel/drv/st.conf  
driver configuration file

/usr/include/sys/mtio.h  
structures and definitions for mag tape io control commands

/usr/include/sys/scsi/targets/stdef.h  
definitions for SCSI tape drives

/dev/rmt/[0-127][l,m,h,u,c][b][n]  
where l, m, h, u, c specifies the density (low, medium, high, ultra/compressed), b the optional BSD behavior (see [mtio\(7I\)](#)), and n the optional no rewind behavior. For example, /dev/rmt/01bn specifies unit 0, low density, BSD behavior, and no rewind.

For 1/2" reel tape devices (HP-88780), the densities are:

|   |                                |
|---|--------------------------------|
| l | 800 BPI density                |
| m | 1600 BPI density               |
| h | 6250 BPI density               |
| c | data compression               |
|   | (not supported on all modules) |

For 8mm tape devices (Exabyte 8200/8500/8505):

|   |                                  |
|---|----------------------------------|
| l | Standard 2 Gbyte format          |
| m | 5 Gbyte format (8500, 8505 only) |

st(7D)

---

|      |                                       |
|------|---------------------------------------|
| h, c | 5 Gbyte compressed format (8505 only) |
|------|---------------------------------------|

---

For 4mm DAT tape devices (Archive Python):

---

|         |                  |
|---------|------------------|
| l       | Standard format  |
| m, h, c | data compression |

---

For all QIC (other than QIC-24) tape devices:

---

|            |                                                                                           |
|------------|-------------------------------------------------------------------------------------------|
| l, m, h, c | density of the tape cartridge type<br>(not all devices can read and<br>write all formats) |
|------------|-------------------------------------------------------------------------------------------|

---

For QIC-24 tape devices (Emulex MT-02):

---

|         |               |
|---------|---------------|
| l       | QIC-11 Format |
| m, h, c | QIC-24 Format |

---

**SEE ALSO** [mt\(1\)](#), [modload\(1M\)](#), [modunload\(1M\)](#), [open\(2\)](#), [read\(2\)](#), [write\(2\)](#), [aioread\(3AIO\)](#), [aiowrite\(3AIO\)](#), [kstat\(3KSTAT\)](#), [driver.conf\(4\)](#), [scsi\(4\)](#), [standards\(5\)](#), [esp\(7D\)](#), [isp\(7D\)](#), [mtio\(7I\)](#), [ioctl\(9E\)](#)

**DIAGNOSTICS** The `st` driver diagnostics may be printed to the console or messages file.

Each diagnostic is dependent on the value of the system variable `st_error_level`. `st_error_level` may be set in the `/etc/system` file. The default setting for `st_error_level` is 4 (`SCSI_ERR_RETRYABLE`) which is suitable for most configurations since only actual fault diagnostics are printed. Settings range from values 0 (`SCSI_ERR_ALL`) which is most verbose, to 6 (`SCSI_ERR_NONE`) which is least verbose. See `stdef.h` for the full list of error-levels. `SCSI_ERR_ALL` level the amount of diagnostic information is likely to be excessive and unnecessary.

The `st` driver diagnostics are described below:

```
Error for Command: <scsi_cmd_name()> Error Level:<error_class>
Requested Block: <blkno> Error Block: <err_blkno>
Vendor: <name>: Serial Number: <inq_serial>
Sense Key: <es_key> ASC: 0x<es_add_code> (scsi_asc_ascq_name()), ASCQ:
0x<es_qual_code>, FRU: 0x<ex_fru_code>
```

where `<error_class>` may be any one of the following: "All," "Unknown," "Informational," "Recovered," "Retryable," "Fatal"

The command indicated by `<scsi_cmd_name>` failed. Requested Block represents the block where the transfer started. Error Block represents the block that caused the error. Sense Key, ASC, ASCQ and FRU information is returned by the target in response to a request sense command. See SCSI protocol documentation for description of Sense Key, ASC, ASCQ, FRU.

The `st` driver attempts to validate entries in the `st.conf` file. Each field in the entry is checked for upper and lower limits and invalid bits set. The fields are named as follows in config string order:

```

conf version
drive type
block size
options
number of densities
density code
default density
non motion timeout
I/O timeout
space timeout
load timeout
unload timeout
erase timeout

```

The `st.conf` diagnostics are described below:

```
<con-name> <field-in-err> <problem-with-field>
```

where `<con-name>` is the name of the config string. Where `<field-in-err>` is the field containing invalid entries and where `<problem-with-field>` describes the nature of the invalid entry.

```
Write/read: not modulo <n> block size
```

The request size for fixed record size devices must be a multiple of the specified block size.

```
Recovery by resets failed
```

After a transport error, the driver attempted to recover by issuing a device reset and then a bus reset if device reset failed. These recoveries failed.

```
Periodic head cleaning required
```

The driver reported that periodic head cleaning is now required. This diagnostic is generated either due to a threshold number of retries, or due to the device communicating to the driver that head cleaning is required.

```
Soft error rate (<n>%) during writing/reading was too high
```

The soft error rate has exceeded the threshold specified by the vendor.

```
SCSI transport failed: reason 'xxxx': {retrying|giving up}
```

The Host Bus Adapter (HBA) has failed to transport a command to the target for the reason stated. The driver will either retry the command or, ultimately, give up.

st(7D)

Tape not inserted in drive

A media access command was attempted while there was no tape inserted into the specified drive. In this case, the drive returns sense key of DRIVE NOT READY.

Transport rejected

The Host Bus Adapter (HBA) driver is not accepting commands after failing to successfully transport a scsi packet to the target. The actual status received by the `st` driver from the underlying HBA driver was either `TRAN_FATAL_ERROR` or `TRAN_BADPKT`.

Retrying command

The `st` driver failed to complete a command. However the command is retryable and will be retried.

Giving up

The `st` driver has exhausted retries or otherwise is unable to retry the command and so is giving up.

No target struct for st%d

The `st` driver failed to obtain state information because the requested state structure was not allocated. The specified device was probably not attached.

File mark detected

The operation detected an end of file mark. (File marks signify the end of a file on the tape media).

End-of-media detected

The operation reached the end of the tape media.

Exabyte soft error reporting failed. DAT soft error reporting failed

The `st` driver was unable to determine if the soft error threshold had been exceeded because it did not successfully read the data it requires or did not obtain enough data. This data is retrieved using the log sense command.

Log sense parameter code does not make sense

The log sense command retrieves hardware statistics that are stored on the drive (for example, soft error counts and retries.) If the data retrieved from the drive is invalid, this message is printed and the data is not used.

Restoring tape position at fileno=%x, blkno=%lx...

The `st` driver is positioning to the specified file and block. This occurs on an open.

Failed to restore the last <file/block> position:

In this state, tape will be loaded at BOT during next open



The `st` driver could not position to the specified location and will revert to the beginning of the tape when the next open is attempted.

```
Device does not support compression
```

The compression facility of the device was requested. However the device does not have a hardware compression capability.

```
DAT soft error reset failed
```

After DAT soft error reporting, the counters within the device that accumulate this sense data need to be re-set. This operation failed.

```
Errors after pkt alloc (b_flags=0x%x, b_error=0x%x)
```

Memory allocation for a scsi packet failed.

```
Incorrect length indicator set
```

The drive reported the length of data requested in a READ operation, was incorrect. Incorrect Length Indicator (ILI) is a very commonly used facility in SCSI tape protocol and should not be seen as an error per-se. Applications typically probe a new tape with a read of any length, using the returned length to the read system call for future reads. Along with this operation, an underlying ILI error is received. ILI errors are therefore informational only and are masked at the default `st_error_level`.

```
Data property (%s) has no value
```

```
Data property (%s) incomplete
```

```
Version # for data property (%s) greater than 1
```

These diagnostics indicate problems in retrieving the values of the various property settings. The `st` driver is in the process of setting the property/parameter values for the tape drive using information from either the built-in table within the driver or from uncommented entries in the `st.conf` file. The effect on the system may be that the tape drive may be set with default or generic driver settings which may not be appropriate for the actual type of tape drive being used.

```
st_attach-RESUME: tape failure tape position will be lost
```

On a resume after a power management suspend, the previously known tape position is no longer valid. This can occur if the tape was changed while the system was in power management suspend. The operation will not be retried.

Write Data Buffering has been deprecated. Your applications should continue to work normally. However, they should be ported to use Asynchronous I/O.

Indicates that buffering has been removed from Solaris.

```
Cannot detach: fileno=%x, blkno=%lx
```

The `st` driver cannot unload because the tape is not positioned at BOT (beginning of tape). May indicate hardware problems with the tape drive.

```
Variable record length I/O
```

```
Fixed record length (%d byte blocks) I/O
```

st(7D)

Tape-drives can use either Fixed or Variable record length. If the drive uses Fixed length records, then the built in property table or the st.conf file will contain a non-zero record-length property. Most DAT, Exabyte and DLT drives support Variable record lengths. Many QIC format tape drives have historically been of Fixed record length.

```
Command will be retried
un_ncmds: %d can't retry cmd
```

These diagnostics are only seen with tape drives with the ST\_RETRY\_ON\_RECOVERED\_DEFERRED\_ERROR bit set. See stdef.h for explanation of the specific usage of this setting.

#### WARNINGS

Effective with Solaris 2.4, the ST\_NO\_RECSIZE\_LIMIT flag is set for the built-in config entries of the Archive DAT and Exabyte drivers by default. (Refer to Large Record Sizes.) Tapes written with large block sizes prior to Solaris 2.4 may cause some applications to fail if the number of bytes returned by a read request is less than the requested block size (for example, asking for 128 Kbytes and receiving less than 64 Kbytes).

The ST\_NO\_RECSIZE\_LIMIT flag can be disabled in the config entry for the device as a work-around. (Refer to Tape Configuration.) This action disables the ability to read and write with large block sizes and allows the reading of tapes written prior to Solaris 2.4 with large block sizes.

(Refer to [mtio\(7I\)](#) for a description of maximum record sizes.)

#### BUGS

Tape devices that do not return a BUSY status during tape loading prevent user commands from being held until the device is ready. The user must delay issuing any tape operations until the tape device is ready. This is not a problem for tape devices supplied by Sun Microsystems.

Tape devices that do not report a blank check error at the end of recorded media may cause file positioning operations to fail. Some tape drives, for example, mistakenly report media error instead of blank check error.

stp4020(7D)

|                    |                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | stp4020 – STP 4020 PCMCIA Adapter                                                                                                                                                                                                                                                                                                                         |
| <b>DESCRIPTION</b> | <p>The STP 4020 PCMCIA Adapter provides for two PCMCIA PC Card sockets. The stp4020 adapter driver provides an interface between the PCMCIA sockets and the PCMCIA nexus. The driver supports the Sun PCMCIA Interface/Sbus card.</p> <p>Direct access to the PCMCIA hardware is not supported. The driver exists solely to support the PCMCIA nexus.</p> |
| <b>FILES</b>       | /kernel/drv/stp4020      stp4020 driver.                                                                                                                                                                                                                                                                                                                  |
| <b>SEE ALSO</b>    | pcmcia(4)                                                                                                                                                                                                                                                                                                                                                 |

streamio(7I)

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|----------------------|---------------------|--------------------------------------------------------|--------------------|------------------------------------|--------------------|------------------------------------|----------------------|---------------------------------------------------|
| <b>NAME</b>          | streamio – STREAMS ioctl commands                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <b>SYNOPSIS</b>      | <pre>#include &lt;sys/types.h&gt; #include &lt;stropts.h&gt; #include &lt;sys/conf.h&gt;  int ioctl(int fildes, int command, ... /*arg*/);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <b>DESCRIPTION</b>   | <p>STREAMS (see intro(3)) ioctl commands are a subset of the ioctl(2) commands and perform a variety of control functions on streams.</p> <p>The <i>fildes</i> argument is an open file descriptor that refers to a stream. The <i>command</i> argument determines the control function to be performed as described below. The <i>arg</i> argument represents additional information that is needed by this command. The type of <i>arg</i> depends upon the command, but it is generally an integer or a pointer to a command-specific data structure. The <i>command</i> and <i>arg</i> arguments are interpreted by the STREAM head. Certain combinations of these arguments may be passed to a module or driver in the stream.</p> <p>Since these STREAMS commands are ioctls, they are subject to the errors described in ioctl(2). In addition to those errors, the call will fail with <code>errno</code> set to <code>EINVAL</code>, without processing a control function, if the STREAM referenced by <i>fildes</i> is linked below a multiplexor, or if <i>command</i> is not a valid value for a stream.</p> <p>Also, as described in ioctl(2), STREAMS modules and drivers can detect errors. In this case, the module or driver sends an error message to the STREAM head containing an error value. This causes subsequent calls to fail with <code>errno</code> set to this value.</p> |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <b>IOCTLS</b>        | <p>The following ioctl commands, with error values indicated, are applicable to all STREAMS files:</p> <table><tr><td><code>I_PUSH</code></td><td>Pushes the module whose name is pointed to by <i>arg</i> onto the top of the current stream, just below the STREAM head. If the STREAM is a pipe, the module will be inserted between the stream heads of both ends of the pipe. It then calls the open routine of the newly-pushed module. On failure, <code>errno</code> is set to one of the following values:</td></tr><tr><td><code>EINVAL</code></td><td>Invalid module name.</td></tr><tr><td><code>EFAULT</code></td><td><i>arg</i> points outside the allocated address space.</td></tr><tr><td><code>ENXIO</code></td><td>Open routine of new module failed.</td></tr><tr><td><code>ENXIO</code></td><td>Hangup received on <i>fildes</i>.</td></tr><tr><td><code>ENOTSUP</code></td><td>Pushing a module is not supported on this stream.</td></tr></table>                                                                                                                                                                                                                                                                                                                                                                                                                                | <code>I_PUSH</code> | Pushes the module whose name is pointed to by <i>arg</i> onto the top of the current stream, just below the STREAM head. If the STREAM is a pipe, the module will be inserted between the stream heads of both ends of the pipe. It then calls the open routine of the newly-pushed module. On failure, <code>errno</code> is set to one of the following values: | <code>EINVAL</code> | Invalid module name. | <code>EFAULT</code> | <i>arg</i> points outside the allocated address space. | <code>ENXIO</code> | Open routine of new module failed. | <code>ENXIO</code> | Hangup received on <i>fildes</i> . | <code>ENOTSUP</code> | Pushing a module is not supported on this stream. |
| <code>I_PUSH</code>  | Pushes the module whose name is pointed to by <i>arg</i> onto the top of the current stream, just below the STREAM head. If the STREAM is a pipe, the module will be inserted between the stream heads of both ends of the pipe. It then calls the open routine of the newly-pushed module. On failure, <code>errno</code> is set to one of the following values:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <code>EINVAL</code>  | Invalid module name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <code>EFAULT</code>  | <i>arg</i> points outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <code>ENXIO</code>   | Open routine of new module failed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <code>ENXIO</code>   | Hangup received on <i>fildes</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |
| <code>ENOTSUP</code> | Pushing a module is not supported on this stream.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                     |                                                                                                                                                                                                                                                                                                                                                                   |                     |                      |                     |                                                        |                    |                                    |                    |                                    |                      |                                                   |

|          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| I_POP    | <p>Removes the module just below the STREAM head of the STREAM pointed to by <i>fildev</i>. To remove a module from a pipe requires that the module was pushed on the side it is being removed from. <i>arg</i> should be 0 in an I_POP request. On failure, <i>errno</i> is set to one of the following values:</p> <p>EINVAL            No module present in the stream.</p> <p>ENXIO            Hangup received on <i>fildev</i>.</p> <p>EPERM            Attempt to pop through an anchor by an unprivileged process.</p> <p>ENOTSUP         Removal is not supported.</p>                                                                          |
| I_ANCHOR | <p>Positions the stream anchor to be at the STREAMS module directly below the STREAM head. Once this has been done, only a privileged process may pop modules below the anchor on the stream. <i>arg</i> must be 0 in an I_ANCHOR request. On failure, <i>errno</i> is set to the following value:</p> <p>EINVAL            Request to put an anchor on a pipe.</p>                                                                                                                                                                                                                                                                                     |
| I_LOOK   | <p>Retrieves the name of the module just below the STREAM head of the STREAM pointed to by <i>fildev</i>, and places it in a null terminated character string pointed at by <i>arg</i>. The buffer pointed to by <i>arg</i> should be at least FMNAMESZ+1 bytes long. This requires the declaration <code>#include &lt;sys/conf.h&gt;</code>. On failure, <i>errno</i> is set to one of the following values:</p> <p>EFAULT            <i>arg</i> points outside the allocated address space.</p> <p>EINVAL            No module present in stream.</p>                                                                                                 |
| I_FLUSH  | <p>This request flushes all input and/or output queues, depending on the value of <i>arg</i>. Legal <i>arg</i> values are:</p> <p>FLUSHR            Flush read queues.</p> <p>FLUSHW            Flush write queues.</p> <p>FLUSHRW          Flush read and write queues.</p> <p>If a pipe or FIFO does not have any modules pushed, the read queue of the STREAM head on either end is flushed depending on the value of <i>arg</i>.</p> <p>If FLUSHR is set and <i>fildev</i> is a pipe, the read queue for that end of the pipe is flushed and the write queue for the other end is flushed. If <i>fildev</i> is a FIFO, both queues are flushed.</p> |

If `FLUSHW` is set and *fdes* is a pipe and the other end of the pipe exists, the read queue for the other end of the pipe is flushed and the write queue for this end is flushed. If *fdes* is a FIFO, both queues of the FIFO are flushed.

If `FLUSHRW` is set, all read queues are flushed, that is, the read queue for the FIFO and the read queue on both ends of the pipe are flushed.

Correct flush handling of a pipe or FIFO with modules pushed is achieved via the `pipemod` module. This module should be the first module pushed onto a pipe so that it is at the midpoint of the pipe itself.

On failure, `errno` is set to one of the following values:

`ENOSR`            Unable to allocate buffers for flush message due to insufficient STREAMS memory resources.

`EINVAL`           Invalid *arg* value.

`ENXIO`            Hangup received on *fdes*.

`I_FLUSHBAND`      Flushes a particular band of messages. *arg* points to a `bandinfo` structure that has the following members:

```
unsigned char bi_pri;
int bi_flag;
```

The `bi_flag` field may be one of `FLUSHR`, `FLUSHW`, or `FLUSHRW` as described earlier.

`I_SETSIG`          Informs the STREAM head that the user wishes the kernel to issue the `SIGPOLL` signal (see `signal(3C)`) when a particular event has occurred on the STREAM associated with *fdes*. `I_SETSIG` supports an asynchronous processing capability in STREAMS. The value of *arg* is a bitmask that specifies the events for which the user should be signaled. It is the bitwise OR of any combination of the following constants:

`S_INPUT`           Any message other than an `M_PCPROTO` has arrived on a STREAM head read queue. This event is maintained for compatibility with previous releases. This event is triggered even if the message is of zero length.

`S_RDNORM`          An ordinary (non-priority) message has arrived on a STREAM head read queue. This event is triggered even if the message is of zero length.

|                        |                                                                                                                                                                                                 |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>S_RDBAND</code>  | A priority band message (band > 0) has arrived on a stream head read queue. This event is triggered even if the message is of zero length.                                                      |
| <code>S_HIPRI</code>   | A high priority message is present on the STREAM head read queue. This event is triggered even if the message is of zero length.                                                                |
| <code>S_OUTPUT</code>  | The write queue just below the STREAM head is no longer full. This notifies the user that there is room on the queue for sending (or writing) data downstream.                                  |
| <code>S_WRNORM</code>  | This event is the same as <code>S_OUTPUT</code> .                                                                                                                                               |
| <code>S_WRBAND</code>  | A priority band greater than 0 of a queue downstream exists and is writable. This notifies the user that there is room on the queue for sending (or writing) priority data downstream.          |
| <code>S_MSG</code>     | A STREAMS signal message that contains the SIGPOLL signal has reached the front of the STREAM head read queue.                                                                                  |
| <code>S_ERROR</code>   | An <code>M_ERROR</code> message has reached the STREAM head.                                                                                                                                    |
| <code>S_HANGUP</code>  | An <code>M_HANGUP</code> message has reached the STREAM head.                                                                                                                                   |
| <code>S_BANDURG</code> | When used in conjunction with <code>S_RDBAND</code> , <code>SIGURG</code> is generated instead of <code>SIGPOLL</code> when a priority message reaches the front of the stream head read queue. |

A user process may choose to be signaled only of high priority messages by setting the *arg* bitmask to the value `S_HIPRI`.

Processes that wish to receive SIGPOLL signals must explicitly register to receive them using `I_SETSIG`. If several processes register to receive this signal for the same event on the same stream, each process will be signaled when the event occurs.

If the value of *arg* is zero, the calling process will be unregistered and will not receive further SIGPOLL signals. On failure, *errno* is set to one of the following values:

|                     |                                                                                                                |
|---------------------|----------------------------------------------------------------------------------------------------------------|
| <code>EINVAL</code> | <i>arg</i> value is invalid or <i>arg</i> is zero and process is not registered to receive the SIGPOLL signal. |
|---------------------|----------------------------------------------------------------------------------------------------------------|

streamio(7I)

|          |        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | EAGAIN | Allocation of a data structure to store the signal request failed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| I_GETSIG |        | Returns the events for which the calling process is currently registered to be sent a SIGPOLL signal. The events are returned as a bitmask pointed to by <i>arg</i> , where the events are those specified in the description of I_SETSIG above. On failure, <i>errno</i> is set to one of the following values:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|          | EINVAL | Process not registered to receive the SIGPOLL signal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|          | EFAULT | <i>arg</i> points outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| I_FIND   |        | Compares the names of all modules currently present in the STREAM to the name pointed to by <i>arg</i> , and returns 1 if the named module is present in the stream. It returns 0 if the named module is not present. On failure, <i>errno</i> is set to one of the following values:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|          | EFAULT | <i>arg</i> points outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|          | EINVAL | <i>arg</i> does not contain a valid module name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| I_PEEK   |        | Allows a user to retrieve the information in the first message on the STREAM head read queue without taking the message off the queue. I_PEEK is analogous to <code>getmsg(2)</code> except that it does not remove the message from the queue. <i>arg</i> points to a <code>strpeek</code> structure, which contains the following members:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|          |        | <pre> struct strbuf ctlbuf; struct strbuf databuf; long flags; </pre> <p>The <code>maxlen</code> field in the <code>ctlbuf</code> and <code>databuf</code> <code>strbuf</code> structures (see <code>getmsg(2)</code>) must be set to the number of bytes of control information and/or data information, respectively, to retrieve. <code>flags</code> may be set to <code>RS_HIPRI</code> or 0. If <code>RS_HIPRI</code> is set, I_PEEK will look for a high priority message on the STREAM head read queue. Otherwise, I_PEEK will look for the first message on the STREAM head read queue.</p> <p>I_PEEK returns 1 if a message was retrieved, and returns 0 if no message was found on the STREAM head read queue. It does not wait for a message to arrive. On return, <code>ctlbuf</code> specifies information in the control buffer, <code>databuf</code> specifies information in the data buffer, and <code>flags</code> contains the value <code>RS_HIPRI</code> or 0. On failure, <i>errno</i> is set to the following value:</p> |
|          | EFAULT | <i>arg</i> points, or the buffer area specified in <code>ctlbuf</code> or <code>databuf</code> is, outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



|          |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | EBADMSG   | Queued message to be read is not valid for I_PEEK.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|          | EINVAL    | Illegal value for flags.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| I_SRDOPT |           | Sets the read mode (see <code>read(2)</code> ) using the value of the argument <i>arg</i> . Legal <i>arg</i> values are:                                                                                                                                                                                                                                                                                                                                                                                   |
|          | RNORM     | Byte-stream mode, the default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|          | RMSGD     | Message-discard mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|          | RMSGN     | Message-nondiscard mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|          |           | In addition, the STREAM head's treatment of control messages may be changed by setting the following flags in <i>arg</i> :                                                                                                                                                                                                                                                                                                                                                                                 |
|          | RPROTNORM | Reject <code>read()</code> with EBADMSG if a control message is at the front of the STREAM head read queue.                                                                                                                                                                                                                                                                                                                                                                                                |
|          | RPROTNORM | Deliver the control portion of a message as data when a user issues <code>read()</code> . This is the default behavior.                                                                                                                                                                                                                                                                                                                                                                                    |
|          | RPROTDIS  | Discard the control portion of a message, delivering any data portion, when a user issues a <code>read()</code> .                                                                                                                                                                                                                                                                                                                                                                                          |
|          |           | On failure, <code>errno</code> is set to the following value:                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|          | EINVAL    | <i>arg</i> is not one of the above legal values, or <i>arg</i> is the bitwise inclusive OR of RMSGD and RMSGN.                                                                                                                                                                                                                                                                                                                                                                                             |
| I_GRDOPT |           | Returns the current read mode setting in an <code>int</code> pointed to by the argument <i>arg</i> . Read modes are described in <code>read()</code> . On failure, <code>errno</code> is set to the following value:                                                                                                                                                                                                                                                                                       |
|          | EFAULT    | <i>arg</i> points outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| I_NREAD  |           | Counts the number of data bytes in data blocks in the first message on the STREAM head read queue, and places this value in the location pointed to by <i>arg</i> . The return value for the command is the number of messages on the STREAM head read queue. For example, if zero is returned in <i>arg</i> , but the <code>ioctl</code> return value is greater than zero, this indicates that a zero-length message is next on the queue. On failure, <code>errno</code> is set to the following value: |
|          | EFAULT    | <i>arg</i> points outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

`I_FDINSERT` Creates a message from specified buffer(s), adds information about another STREAM and sends the message downstream. The message contains a control part and an optional data part. The data and control parts to be sent are distinguished by placement in separate buffers, as described below.

The *arg* argument points to a `strfdinsert` structure, which contains the following members:

```
struct strbuf ctlbuf;
struct strbuf databuf;
t_uscalar_t flags;
int fildes;
int offset;
```

The `len` member in the `ctlbuf` `strbuf` structure (see `putmsg(2)`) must be set to the size of a `t_uscalar_t` plus the number of bytes of control information to be sent with the message. The `fildes` member specifies the file descriptor of the other STREAM, and the `offset` member, which must be suitably aligned for use as a `t_uscalar_t`, specifies the offset from the start of the control buffer where `I_FDINSERT` will store a `t_uscalar_t` whose interpretation is specific to the STREAM end. The `len` member in the `databuf` `strbuf` structure must be set to the number of bytes of data information to be sent with the message, or to 0 if no data part is to be sent.

The `flags` member specifies the type of message to be created. A normal message is created if `flags` is set to 0, and a high-priority message is created if `flags` is set to `RS_HIPRI`. For non-priority messages, `I_FDINSERT` will block if the STREAM write queue is full due to internal flow control conditions. For priority messages, `I_FDINSERT` does not block on this condition. For non-priority messages, `I_FDINSERT` does not block when the write queue is full and `O_NDELAY` or `O_NONBLOCK` is set. Instead, it fails and sets `errno` to `EAGAIN`.

`I_FDINSERT` also blocks, unless prevented by lack of internal resources, waiting for the availability of message blocks in the STREAM, regardless of priority or whether `O_NDELAY` or `O_NONBLOCK` has been specified. No partial message is sent.

The `ioctl()` function with the `I_FDINSERT` command will fail if:

|                     |                                                                                                                                                                                    |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>EAGAIN</code> | A non-priority message is specified, the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flag is set, and the STREAM write queue is full due to internal flow control conditions. |
| <code>ENOSR</code>  | Buffers can not be allocated for the message that is to be created.                                                                                                                |

|        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EFAULT | The <i>arg</i> argument points, or the buffer area specified in <i>ctlbuf</i> or <i>databuf</i> is, outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                    |
| EINVAL | One of the following: The <i>fildev</i> member of the <i>strfdinsert</i> structure is not a valid, open STREAM file descriptor; the size of a <i>t_uscalar_t</i> plus <i>offset</i> is greater than the <i>len</i> member for the buffer specified through <i>ctlptr</i> ; the <i>offset</i> member does not specify a properly-aligned location in the data buffer; or an undefined value is stored in <i>flags</i> .                                                                      |
| ENXIO  | Hangup received on the <i>fildev</i> argument of the <i>ioctl</i> call or the <i>fildev</i> member of the <i>strfdinsert</i> structure.                                                                                                                                                                                                                                                                                                                                                     |
| ERANGE | The <i>len</i> field for the buffer specified through <i>databuf</i> does not fall within the range specified by the maximum and minimum packet sizes of the topmost STREAM module; or the <i>len</i> member for the buffer specified through <i>databuf</i> is larger than the maximum configured size of the data part of a message; or the <i>len</i> member for the buffer specified through <i>ctlbuf</i> is larger than the maximum configured size of the control part of a message. |

*I\_FDINSERT* can also fail if an error message was received by the STREAM head of the STREAM corresponding to the *fildev* member of the *strfdinsert* structure. In this case, *errno* will be set to the value in the message.

**I\_STR**

Constructs an internal STREAMS *ioctl* message from the data pointed to by *arg*, and sends that message downstream.

This mechanism is provided to send user *ioctl* requests to downstream modules and drivers. It allows information to be sent with the *ioctl*, and will return to the user any information sent upstream by the downstream recipient. *I\_STR* blocks until the system responds with either a positive or negative acknowledgement message, or until the request "times out" after some period of time. If the request times out, it fails with *errno* set to *ETIME*.

To send requests downstream, *arg* must point to a *strioc1* structure which contains the following members:

## streamio(7I)

```
int ic_cmd;
int ic_timeout;
int ic_len;
char *ic_dp;
```

`ic_cmd` is the internal `ioctl` command intended for a downstream module or driver and `ic_timeout` is the number of seconds (-1 = infinite, 0 = use default, >0 = as specified) an `I_STR` request will wait for acknowledgement before timing out. `ic_len` is the number of bytes in the data argument and `ic_dp` is a pointer to the data argument. The `ic_len` field has two uses: on input, it contains the length of the data argument passed in, and on return from the command, it contains the number of bytes being returned to the user (the buffer pointed to by `ic_dp` should be large enough to contain the maximum amount of data that any module or the driver in the STREAM can return).

At most one `I_STR` can be active on a stream. Further `I_STR` calls will block until the active `I_STR` completes via a positive or negative acknowledgment, a timeout, or an error condition at the STREAM head. By setting the `ic_timeout` field to 0, the user is requesting STREAMS to provide the "DEFAULT" timeout. The default timeout is specific to the STREAMS implementation and may vary depending on which release of Solaris you are using. For Solaris 8 (and earlier versions), the default timeout is fifteen seconds. The `O_NDELAY` and `O_NONBLOCK` (see `open(2)`) flags have no effect on this call.

The STREAM head will convert the information pointed to by the `strioctl` structure to an internal `ioctl` command message and send it downstream. On failure, `errno` is set to one of the following values:

|                     |                                                                                                                                                                                                                                     |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ENOSR</code>  | Unable to allocate buffers for the <code>ioctl</code> message due to insufficient STREAMS memory resources.                                                                                                                         |
| <code>EFAULT</code> | Either <code>arg</code> points outside the allocated address space, or the buffer area specified by <code>ic_dp</code> and <code>ic_len</code> (separately for data sent and data returned) is outside the allocated address space. |
| <code>EINVAL</code> | <code>ic_len</code> is less than 0 or <code>ic_len</code> is larger than the maximum configured size of the data part of a message or <code>ic_timeout</code> is less than -1.                                                      |
| <code>ENXIO</code>  | Hangup received on <i>files</i> .                                                                                                                                                                                                   |
| <code>ETIME</code>  | A downstream <code>ioctl</code> timed out before acknowledgement was received.                                                                                                                                                      |

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                       | An <code>I_STR</code> can also fail while waiting for an acknowledgement if a message indicating an error or a hangup is received at the <code>STREAM</code> head. In addition, an error code can be returned in the positive or negative acknowledgement message, in the event the <code>ioctl</code> command sent downstream fails. For these cases, <code>I_STR</code> will fail with <code>errno</code> set to the value in the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <code>I_SWROPT</code> | <p>Sets the write mode using the value of the argument <code>arg</code>. Legal bit settings for <code>arg</code> are:</p> <p><code>SNDZERO</code>            Send a zero-length message downstream when a write of 0 bytes occurs.</p> <p>To not send a zero-length message when a write of 0 bytes occurs, this bit must not be set in <code>arg</code>.</p> <p>On failure, <code>errno</code> may be set to the following value:</p> <p><code>EINVAL</code>            <code>arg</code> is not the above legal value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <code>I_GWROPT</code> | Returns the current write mode setting, as described above, in the <code>int</code> that is pointed to by the argument <code>arg</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <code>I_SENDFD</code> | <p>Requests the <code>STREAM</code> associated with <code>fildev</code> to send a message, containing a file pointer, to the stream head at the other end of a <code>STREAM</code> pipe. The file pointer corresponds to <code>arg</code>, which must be an open file descriptor.</p> <p><code>I_SENDFD</code> converts <code>arg</code> into the corresponding system file pointer. It allocates a message block and inserts the file pointer in the block. The user id and group id associated with the sending process are also inserted. This message is placed directly on the read queue (see <code>intro(3)</code>) of the <code>STREAM</code> head at the other end of the <code>STREAM</code> pipe to which it is connected. On failure, <code>errno</code> is set to one of the following values:</p> <p><code>EAGAIN</code>            The sending <code>STREAM</code> is unable to allocate a message block to contain the file pointer.</p> <p><code>EAGAIN</code>            The read queue of the receiving <code>STREAM</code> head is full and cannot accept the message sent by <code>I_SENDFD</code>.</p> <p><code>EBADF</code>            <code>arg</code> is not a valid, open file descriptor.</p> <p><code>EINVAL</code>            <code>fildev</code> is not connected to a <code>STREAM</code> pipe.</p> <p><code>ENXIO</code>            Hangup received on <code>fildev</code>.</p> |
| <code>I_RECVFD</code> | Retrieves the file descriptor associated with the message sent by an <code>I_SENDFD</code> <code>ioctl</code> over a <code>STREAM</code> pipe. <code>arg</code> is a pointer to a data buffer large enough to hold an <code>strrecvfd</code> data structure containing the following members:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## streamio(7I)

```
int fd;
uid_t uid;
gid_t gid;
```

`fd` is an integer file descriptor. `uid` and `gid` are the user id and group id, respectively, of the sending stream.

If `O_NDELAY` and `O_NONBLOCK` are clear (see `open(2)`), `I_RECVFD` will block until a message is present at the STREAM head. If `O_NDELAY` or `O_NONBLOCK` is set, `I_RECVFD` will fail with `errno` set to `EAGAIN` if no message is present at the STREAM head.

If the message at the STREAM head is a message sent by an `I_SENDFD`, a new user file descriptor is allocated for the file pointer contained in the message. The new file descriptor is placed in the `fd` field of the `strrecvfd` structure. The structure is copied into the user data buffer pointed to by `arg`. On failure, `errno` is set to one of the following values:

|                         |                                                                                                                               |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <code>EAGAIN</code>     | A message is not present at the STREAM head read queue, and the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flag is set. |
| <code>EBADMSG</code>    | The message at the STREAM head read queue is not a message containing a passed file descriptor.                               |
| <code>EFAULT</code>     | <code>arg</code> points outside the allocated address space.                                                                  |
| <code>EMFILE</code>     | <code>NOFILES</code> file descriptors are currently open.                                                                     |
| <code>ENXIO</code>      | Hangup received on <i>files</i> .                                                                                             |
| <code>E_OVERFLOW</code> | <code>uid</code> or <code>gid</code> is too large to be stored in the structure pointed to by <code>arg</code> .              |

### `I_LIST`

Allows the user to list all the module names on the stream, up to and including the topmost driver name. If `arg` is `NULL`, the return value is the number of modules, including the driver, that are on the STREAM pointed to by *files*. This allows the user to allocate enough space for the module names. If `arg` is non-null, it should point to an `str_list` structure that has the following members:

```
int sl_nmods;
struct str_mlist *sl_modlist;
```

The `str_mlist` structure has the following member:

```
char l_name[FMNAMESZ+1];
```

The `sl_nmods` member indicates the number of entries the process has allocated in the array. Upon return, the `sl_modlist` member of the `str_list` structure contains the list of module names, and the number of entries that have been filled into the `sl_modlist` array is found in the `sl_nmods` member (the

number includes the number of modules including the driver). The return value from `ioctl()` is 0. The entries are filled in starting at the top of the STREAM and continuing downstream until either the end of the STREAM is reached, or the number of requested modules (`s1_nmods`) is satisfied. On failure, `errno` may be set to one of the following values:

EINVAL           The `s1_nmods` member is less than 1.

EAGAIN           Unable to allocate buffers

**I\_ATMARK**       Allows the user to see if the current message on the stream head read queue is “marked” by some module downstream. *arg* determines how the checking is done when there may be multiple marked messages on the STREAM head read queue. It may take the following values:

ANYMARK          Check if the message is marked.

LASTMARK         Check if the message is the last one marked on the queue.

The return value is 1 if the mark condition is satisfied and 0 otherwise. On failure, `errno` is set to the following value:

EINVAL           Invalid *arg* value.

**I\_CKBAND**       Check if the message of a given priority band exists on the stream head read queue. This returns 1 if a message of a given priority exists, 0 if not, or -1 on error. *arg* should be an integer containing the value of the priority band in question. On failure, `errno` is set to the following value:

EINVAL           Invalid *arg* value.

**I\_GETBAND**      Returns the priority band of the first message on the STREAM head read queue in the integer referenced by *arg*. On failure, `errno` is set to the following value:

ENODATA          No message on the STREAM head read queue.

**I\_CANPUT**       Check if a certain band is writable. *arg* is set to the priority band in question. The return value is 0 if the priority band *arg* is flow controlled, 1 if the band is writable, or -1 on error. On failure, `errno` is set to the following value:

EINVAL           Invalid *arg* value.

**I\_SETCLTIME**   Allows the user to set the time the STREAM head will delay when a stream is closing and there are data on the write queues. Before closing each module and driver, the STREAM head will delay for the specified amount of time to allow the data to drain. Note, however, that the module or driver may itself delay in its close

|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|             | routine; this delay is independent of the STREAM head's delay and is not settable. If, after the delay, data are still present, data will be flushed. <i>arg</i> is a pointer to an integer containing the number of milliseconds to delay, rounded up to the nearest legal value on the system. The default is fifteen seconds. On failure, <i>errno</i> is set to the following value:                                                                                                                                       |
|             | EINVAL                    Invalid <i>arg</i> value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| I_GETCLTIME | Returns the close time delay in the integer pointed by <i>arg</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| I_SERROPT   | Sets the error mode using the value of the argument <i>arg</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|             | Normally STREAM head errors are persistent; once they are set due to an M_ERROR or M_HANGUP, the error condition will remain until the STREAM is closed. This option can be used to set the STREAM head into non-persistent error mode i.e. once the error has been returned in response to a read(2), getmsg(2), ioctl(2), write(2), or putmsg(2) call the error condition will be cleared. The error mode can be controlled independently for read and write side errors. Legal <i>arg</i> values are either none or one of: |
|             | RERRNORM                    Persistent read errors, the default.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|             | RERRNONPERSIST            Non-persistent read errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|             | OR'ed with either none or one of:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|             | WERRNORM                    Persistent write errors, the default.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|             | WERRNONPERSIST            Non-persistent write errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|             | When no value is specified e.g. for the read side error behavior then the behavior for that side will be left unchanged.                                                                                                                                                                                                                                                                                                                                                                                                       |
|             | On failure, <i>errno</i> is set to the following value:                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|             | EINVAL <i>arg</i> is not one of the above legal values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| I_GERROPT   | Returns the current error mode setting in an <i>int</i> pointed to by the argument <i>arg</i> . Error modes are described above for I_SERROPT. On failure, <i>errno</i> is set to the following value:                                                                                                                                                                                                                                                                                                                         |
|             | EFAULT <i>arg</i> points outside the allocated address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|             | The following four commands are used for connecting and disconnecting multiplexed STREAMS configurations.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| I_LINK      | Connects two streams, where <i>fildev</i> is the file descriptor of the stream connected to the multiplexing driver, and <i>arg</i> is the file descriptor of the STREAM connected to another driver. The                                                                                                                                                                                                                                                                                                                      |



STREAM designated by *arg* gets connected below the multiplexing driver. `I_LINK` requires the multiplexing driver to send an acknowledgement message to the STREAM head regarding the linking operation. This call returns a multiplexor ID number (an identifier used to disconnect the multiplexor, see `I_UNLINK`) on success, and -1 on failure. On failure, `errno` is set to one of the following values:

|        |                                                                                                                                                                                  |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENXIO  | Hangup received on <i>fildev</i> .                                                                                                                                               |
| ETIME  | Time out before acknowledgement message was received at STREAM head.                                                                                                             |
| EAGAIN | Temporarily unable to allocate storage to perform the <code>I_LINK</code> .                                                                                                      |
| ENOSR  | Unable to allocate storage to perform the <code>I_LINK</code> due to insufficient STREAMS memory resources.                                                                      |
| EBADF  | <i>arg</i> is not a valid, open file descriptor.                                                                                                                                 |
| EINVAL | <i>fildev</i> STREAM does not support multiplexing.                                                                                                                              |
| EINVAL | <i>arg</i> is not a stream, or is already linked under a multiplexor.                                                                                                            |
| EINVAL | The specified link operation would cause a "cycle" in the resulting configuration; that is, a driver would be linked into the multiplexing configuration in more than one place. |
| EINVAL | <i>fildev</i> is the file descriptor of a pipe or FIFO.                                                                                                                          |
| EINVAL | Either the upper or lower stream has a major number $\geq$ the maximum major number on the system.                                                                               |

An `I_LINK` can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the STREAM head of *fildev*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, `I_LINK` will fail with `errno` set to the value in the message.

`I_UNLINK`

Disconnects the two streams specified by *fildev* and *arg*. *fildev* is the file descriptor of the STREAM connected to the multiplexing driver. *arg* is the multiplexor ID number that was returned by the `I_LINK`. If *arg* is -1, then all streams that were linked to *fildev* are disconnected. As in `I_LINK`, this command requires the multiplexing driver to acknowledge the unlink. On failure, `errno` is set to one of the following values:

streamio(7I)

|        |                                                                                                                                               |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| ENXIO  | Hangup received on <i>fildev</i> .                                                                                                            |
| ETIME  | Time out before acknowledgement message was received at STREAM head.                                                                          |
| ENOSR  | Unable to allocate storage to perform the I_UNLINK due to insufficient STREAMS memory resources.                                              |
| EINVAL | <i>arg</i> is an invalid multiplexor ID number or <i>fildev</i> is not the STREAM on which the I_LINK that returned <i>arg</i> was performed. |
| EINVAL | <i>fildev</i> is the file descriptor of a pipe or FIFO.                                                                                       |

An I\_UNLINK can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the STREAM head of *fildev*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, I\_UNLINK will fail with *errno* set to the value in the message.

I\_PLINK

Connects two streams, where *fildev* is the file descriptor of the stream connected to the multiplexing driver, and *arg* is the file descriptor of the STREAM connected to another driver. The STREAM designated by *arg* gets connected via a persistent link below the multiplexing driver. I\_PLINK requires the multiplexing driver to send an acknowledgement message to the STREAM head regarding the linking operation. This call creates a persistent link that continues to exist even if the file descriptor *fildev* associated with the upper STREAM to the multiplexing driver is closed. This call returns a multiplexor ID number (an identifier that may be used to disconnect the multiplexor, see I\_PUNLINK) on success, and -1 on failure. On failure, *errno* is set to one of the following values:

|        |                                                                          |
|--------|--------------------------------------------------------------------------|
| ENXIO  | Hangup received on <i>fildev</i> .                                       |
| ETIME  | Time out before acknowledgement message was received at the STREAM head. |
| EAGAIN | Unable to allocate STREAMS storage to perform the I_PLINK.               |
| EBADF  | <i>arg</i> is not a valid, open file descriptor.                         |
| EINVAL | <i>fildev</i> does not support multiplexing.                             |
| EINVAL | <i>arg</i> is not a STREAM or is already linked under a multiplexor.     |

EINVAL The specified link operation would cause a “cycle” in the resulting configuration; that is, if a driver would be linked into the multiplexing configuration in more than one place.

EINVAL *fildev* is the file descriptor of a pipe or FIFO.

An I\_PLINK can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error on a hangup is received at the STREAM head of *fildev*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, I\_PLINK will fail with `errno` set to the value in the message.

I\_PUNLINK Disconnects the two streams specified by *fildev* and *arg* that are connected with a persistent link. *fildev* is the file descriptor of the STREAM connected to the multiplexing driver. *arg* is the multiplexor ID number that was returned by I\_PLINK when a STREAM was linked below the multiplexing driver. If *arg* is MUXID\_ALL then all streams that are persistent links to *fildev* are disconnected. As in I\_PLINK, this command requires the multiplexing driver to acknowledge the unlink. On failure, `errno` is set to one of the following values:

ENXIO Hangup received on *fildev*.

ETIME Time out before acknowledgement message was received at the STREAM head.

EAGAIN Unable to allocate buffers for the acknowledgement message.

EINVAL Invalid multiplexor ID number.

EINVAL *fildev* is the file descriptor of a pipe or FIFO.

An I\_PUNLINK can also fail while waiting for the multiplexing driver to acknowledge the link request if a message indicating an error or a hangup is received at the STREAM head of *fildev*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, I\_PUNLINK will fail with `errno` set to the value in the message.

**RETURN VALUES** Unless specified otherwise above, the return value from `ioctl()` is 0 upon success and -1 upon failure, with `errno` set as indicated.

**SEE ALSO** `intro(3)`, `close(2)`, `fcntl(2)`, `getmsg(2)`, `ioctl(2)`, `open(2)`, `poll(2)`, `putmsg(2)`, `read(2)`, `write(2)`, `signal(3C)`, `signal.h(3HEAD)`,

*STREAMS Programming Guide*

su(7D)

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                              | su – asynchronous serial port driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>SYNOPSIS</b>                          | <pre>#include &lt;fcntl.h&gt; #include &lt;sys/termios.h&gt; open("/dev/tty[a-z]", _mode); open("/dev/term[a-z]", _mode); open("/dev/cua[a-z]", _mode);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>DESCRIPTION</b>                       | <p>The su module is a loadable STREAMS driver that provides basic support for standard UARTS that use Intel-8250, National Semiconductor-16450/16550 hardware and Southbridge 1535D (16550 compatible) Super I/O hardware. The module also provides keyboard and mouse I/O support for Sun machines using those same Intel, National Semiconductor and Southbridge chipsets. The su driver provides basic asynchronous communication support for serial ports. Both the serial devices and keyboard/mouse devices will have streams built with appropriate modules pushed atop the su driver by means of either the <code>autopush(1M)</code> or <code>dacf.conf(4)</code> facilities, depending on the OS revision and architecture in use.</p> <p>The su module supports those <a href="#">termio(7I)</a> device control functions specified by flags in the <code>c_cflag</code> word of the <code>termios</code> structure, and by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code>, or <code>INPCK</code> flags in the <code>c_iflag</code> word of the <code>termios</code> structure. All other <a href="#">termio(7I)</a> functions must be performed by STREAMS modules pushed atop the driver. When a device is opened, the <a href="#">ldterm(7M)</a> and <a href="#">ttcompat(7M)</a> STREAMS modules are automatically pushed on top of the stream, providing the standard <a href="#">termio(7I)</a> interface.</p> <p>The character-special devices <code>/dev/ttya</code> and <code>/dev/ttyb</code> are used to access the two standard serial ports. The su module supports up to ten serial ports, including the standard ports. The <code>tty[a-z]</code> devices have minor device numbers in the range 00-03, and may be assigned names of the form <code>/dev/ttyd_n_</code>, where <code>_n_</code> denotes the line to be accessed. These device names are typically used to provide a logical access point for a <i>dial-in</i> line that is used with a modem.</p> <p>To allow a single tty line to be connected to a modem and used for incoming and outgoing calls, a special feature is available that is controlled by the minor device number. By accessing character-special devices with names of the form <code>/dev/cua_n_</code>, it is possible to open a port without the Carrier Detect signal being asserted, either through hardware or an equivalent software mechanism. These devices are commonly known as <i>dial-out</i> lines.</p> |
| <b>APPLICATION PROGRAMMING INTERFACE</b> | <p>Once a <code>/dev/cua_n_</code> line is opened, the corresponding tty, or ttyd line cannot be opened until the <code>/dev/cua_n_</code> line is closed. A blocking open will wait until the <code>/dev/cua_n_</code> line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again. A non-blocking open will return an error. If the <code>/dev/ttyd_n_</code> line has been opened successfully (usually only when carrier is recognized on the modem), the</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

corresponding `/dev/cua_n_` line cannot be opened. This allows a modem to be attached to a device, (for example, `/dev/ttyd0`, which is renamed from `/dev/tty00`) and used for dial-in (by enabling the line for login in `/etc/inittab`) or dial-out (by `tip(1)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line.

**IOCTLS** The standard set of `termio ioctl()` calls are supported by `su`.

Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls.

The input and output line speeds may be set to any of the following baud rates: 0, 50, 75, 110, 134, 150, 200, 300, 600, 1200, 1800, 2400, 4800, 9600, 19200, 38400, 57600 or 115200. The speeds cannot be set independently; for example, when the output speed is set, the input speed is automatically set to the same speed.

When the `su` module is used to service the serial console port, it supports a BREAK condition that allows the system to enter the debugger or the monitor. The BREAK condition is generated by hardware and it is usually enabled by default.

A BREAK condition originating from erroneous electrical signals cannot be distinguished from one deliberately sent by remote DCE. The Alternate Break sequence can be used as a remedy against this. Due to a risk of incorrect sequence interpretation, SLIP and certain other binary protocols should not be run over the serial console port when Alternate Break sequence is in effect. Although PPP is a binary protocol, it is able to avoid these sequences using the ACCM feature in *RFC 1662*. For Solaris PPP 4.0, you do this by adding the following line to the `/etc/ppp/options` file (or other configuration files used for the connection; see `pppd(1M)` for details):

```
asynctmap 0x00002000
```

By default, the Alternate Break sequence is a three character sequence: carriage return, tilde and control-B (CR ~ CTRL-B), but may be changed by the driver. For more information on breaking (entering the debugger or monitor), see `kbd(1)` and `kb(7M)`.

**ERRORS** An `open()` will fail under the following conditions:

|       |                                                                                                                                                                                   |
|-------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENXIO | The unit being opened does not exist.                                                                                                                                             |
| EBUSY | The dial-out device is being opened while the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open. |
| EBUSY | The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call.                                                                           |

|              |                              |                                           |
|--------------|------------------------------|-------------------------------------------|
| <b>FILES</b> | <code>/dev/cua/[a-z]</code>  | dial-out tty lines                        |
|              | <code>/dev/term/[a-z]</code> | dial-in tty lines                         |
|              | <code>/dev/tty[a-z]</code>   | binary compatibility package device names |

su(7D)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Architecture   | SPARC           |

**SEE ALSO** `strconf(1)`, `kbd(1)`, `tip(1)`, `uucp(1C)`, `autopush(1M)`, `kstat(1M)`, `pppd(1M)`, `ioctl(2)`, `open(2)`, `termios(3C)`, `dacf.conf(4)`, `attributes(5)`, `kb(7M)`, `ldterm(7M)`, `ttcompat(7M)`, `termio(7I)`

**DIAGNOSTICS** The `su` driver keeps track of various warning and error conditions using `kstat` counters. The output of the `kstat su` command provides `kstat` counters. The counters and their meaning follow:

|                      |                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| silo overflow        | The internal chip FIFO received more data than it could handle. This indicates that the Solaris operating environment was not servicing data interrupts fast enough possibly due to a system with too many interrupts or a data line with a data rate that is too high.                                                                                                              |
| ring buffer overflow | The <code>su</code> module was unable to store data it removed from the chips internal FIFO into a software buffer. The user process is not reading data fast enough, possibly due to an overloaded system. If possible, the application should enable flow control (either CTSRTS or XONXOFF) to allow the driver to backpressure the remote system when the local buffers fill up. |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | sxp – Rockwell 2200 SNAP Streams Driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>SYNOPSIS</b>    | <code>/dev/sxp</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>DESCRIPTION</b> | <p>The <code>sxp</code> (also known as the SNAP ) driver is a loadable, clonable, STREAMS driver that supports the connectionless Data Link Provider Interface ( <code>dlpi(7P)</code>) over one or more FDDI adapters (Rockwell 2200 Series). The cloning character-special devices (<code>/dev/sxp</code>, <code>/dev/snap</code>, <code>/dev/llc</code>, <code>/dev/mac</code>) are used to access the 2200 Series adapter(s). The <code>/dev/sxp</code> device is equivalent to <code>/dev/snap</code>. <code>/dev/sxp</code> is used so that the name SXP will show up in <code>ifconfig</code>. All messages transmitted on a SNAP device have the 802.2 LLC and Sub-Network Access Protocol (SNAP) and the FDDI MAC headers ( RFC -1188) prepended. For an LLC device, the LLC and MAC headers are prepended, and for a MAC device only the MAC header is prepended. Received FDDI frames are delivered to the appropriate open device. In response to a <code>DL_INFO_REQ</code>, the SNAP driver returns the following values in the <code>DL_INFO_ACK</code> primitive:</p> <ul style="list-style-type: none"> <li>■ The maximum SDU is 4500.</li> <li>■ The minimum SDU is 0.</li> <li>■ The DLSAP address length is 8 (always true in the Solaris environment).</li> <li>■ The address offset is 0 (prior to being attached).</li> <li>■ The MAC type is <code>DL_FDDI</code>.</li> <li>■ The <code>sap</code> length value is -2, which indicates that within the DLSAP address, the physical address component is followed immediately by a 2-byte service access point ( <code>SAP</code> ) component.</li> <li>■ The service mode is <code>DL_CLDLS</code>.</li> <li>■ The quality of service (QOS) fields are 0, because optional QOS is not supported.</li> <li>■ The provider style is <code>DL_STYLE2</code>.</li> <li>■ The broadcast address value is the IEEE broadcast address (<code>FF:FF:FF:FF:FF:FF</code>)</li> </ul> <p>Because the SNAP driver is a "style 2" Data Link Service provider, an explicit <code>DL_ATTACH_REQ</code> message from the user is required to associate the opened stream with a particular network device (that is, <code>ppa</code>). The <code>dl_ppa</code> field within the <code>DL_ATTACH_REQ</code> indicates the instance (unit) number of the network device. If no currently attached <code>ppa</code> has the same instance number and there are no unattached <code>ppas</code> available, the driver returns an error (<code>DL_ERROR_ACK</code>). Once in the <code>DL_ATTACHED</code> state, a <code>DL_BIND_REQ</code> is required to associate a particular SAP with the stream.</p> <p>Once in the <code>DL_ATTACHED</code> state, a <code>DL_BIND_REQ</code> is required to associate a particular Service Access Point ( <code>SAP</code> ) with the stream. For the <code>sap</code> field within the <code>DL_BIND_REQ</code>, valid values are in the range [0-0xFFFF]. Values for 0-0xFF will give LLC 802.2 service without SNAP encapsulation, unless a later <code>DL_HIERARCHIAL_BIND</code> <code>DL_SUBS_BIND_REQ</code> is made. Values from 0x100-0xFFFF will give LLC 802.2 with SNAP encapsulation without the need for a</p> |

sxp(7D)

DL\_SUBS\_BIND\_REQ. Note that DL\_HIERARCHIAL\_BIND class DL\_SUBS\_BIND\_REQs are only supported on streams bound to the 0xAA SAP. After successful completion of the DL\_BIND\_REQ, the ppa is initialized and the stream is ready for use. In addition to the DL\_HIERARCHIAL\_BIND class of DL\_SUBS\_BUD\_REQ, the DL\_PEER\_BIND class can be used to bind multiple SAP s with a stream.

Frames may be transmitted on the FDDI ring by sending DL\_UNITDATA\_REQ messages to the SNAP driver. The DLSAP address contained within the DL\_UNITDATA\_REQ must consist of both the SAP and physical (FDDI) components. For a SNAP device, the SAP portion of the DLSAP address is placed in the EtherType field of the 802.2 SNAP header. The DSAP and SSAP fields of the 802.2 LLC header are both set to the value 170, indicating a SNAP message and a MAC frame\_type of LLC. For an LLC device, the SAP portion of the DLSAP address is placed in the DSAP field of the 802.2 LLC header. The SSAP field is set to the SAP bound to the stream. The MAC frame\_type is LLC. For a MAC device, the SAP portion of the DLSAP address is placed in the frame\_control field of the MAC header. Received FDDI frames are routed up the correct stream(s) as DL\_UNITDATA\_IND messages (containing the DLSAP address). The stream(s) are found by:

1. Comparing the EtherType field of the SNAP header with the bound SAP of all of the SNAP streams
2. Comparing the DSAP field of the LLC header with the bound SAP of all the LLC streams
3. Comparing the frame\_control field of the MAC header with the bound SAP of all the MAC streams.

If necessary, messages are duplicated. In addition to the mandatory connectionless DLPI message set, the driver also supports the following primitives:

DL\_ENABMULTI\_REQ, DL\_DISABMULTI\_REQ, DL\_PROMISCON\_REQ, DL\_PROMISCOFF\_REQ, DL\_PHYS\_ADDR\_REQ.

The DL\_ENABMULTI\_REQ and DL\_DISABMULTI\_REQ primitives enable or disable reception of individual multicast group addresses. Using these primitives, a set of multicast group addresses may be iteratively created and modified on a per-stream basis. These primitives are accepted by the driver in any state following a successful DL\_ATTACH\_REQ. The DL\_PROMISCON\_REQ and DL\_PROMISCOFF\_REQ primitives (with the DL\_PROMISC\_PHYS flag set in the dl\_levelfield) enable or disable reception of all (promiscuous mode) frames on the media, including frames generated by the local host. When used with the DL\_PROMISC\_SAP flag (set), this enables or disables reception of all sap values. When used with the DL\_PROMISC\_MULTI flag (set), this enables or disables reception of all multicast group addresses. The affect of each primitive is always on a per-stream basis, and is independent of the other sap and physical level configurations on this stream (or other streams). In the DL\_PHYS\_ADDR\_ACK message, the DL\_PHYS\_ADDR\_REQ primitive returns the 6-octet FDDI address (in canonical form) currently associated with the stream. This primitive is valid only in states following a successful DL\_ATTACH\_REQ. The driver also



supports the following *ioctl*s (I/O controls): `DLIOCRAW`, `SL_RAW`, `SL_DATA_ENABLE`, `SL_DATA_DISABLE`, and `DRV_CONFIG`. As defined by Solaris, the `DLIOCRAW ioctl` puts the stream into raw mode, which causes the driver to send the full MAC -level packet up the stream in an `M_DATA` message, instead of transforming it to the `DL_UNITDATA_IND` form. On this stream, the driver will also accept formatted `M_DATA` messages for transmission. To disable raw mode, the stream must be closed. The `DLIOCRAW ioctl` requires no arguments. As defined by Rockwell, the `SL_RAW ioctl` puts the stream into raw mode, similar to the `DLIOCRAW ioctl` except that the frame-type field of the MAC header is considered to be a long word instead of a byte, preserving alignment. The `SL_RAW ioctl` requires no arguments. As defined by Rockwell, the `SL_DATA_ENABLE` and `SL_DATA_DISABLE ioctl`s enable or disable the transmission of data on the stream. By default, transmission is enabled. The `SL_DATA_ENABLE` and `SL_DATA_DISABLE ioctl`s require no arguments.

**FILES** /dev/sxp SXP special character device  
 kernel/drv/sys\_core SXP loadable module  
 kernel/drv/sxp.conf SXP configuration file

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Architecture   | x86             |

**SEE ALSO** `attributes(5)`, `dlpi(7P)`

symhisl(7D)

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>          | symhisl – symhisl SCSI Host Bus Adapter Driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |
| <b>SYNOPSIS</b>      | <code>scsi@unit-address</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |
| <b>DESCRIPTION</b>   | <p>The <code>symhisl</code> host bus adapter driver is a SCSI-compliant nexus driver that supports the LSI Logic SYM21002 (SYM53C896), SYM22910 (SYM53C896), SYM53C895A, SYM53C896, SYM53C1010-33, and SYM53C1010-66 SCSI controller chips.</p> <p>The <code>symhisl</code> driver supports the standard functions provided by the SCSI interface, including tagged and untagged queuing, Wide, Fast, Ultra, Ultra2, and Ultra3 SCSI, and auto request sense. The <code>symhisl</code> driver does not support linked commands.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |
| <b>CONFIGURATION</b> | <p>You configure the <code>symhisl</code> driver by defining properties in the <code>symhisl.conf</code> file. Properties in the <code>symhisl.conf</code> file override the global SCSI settings. The driver supports the following user-modifiable properties:</p> <pre>scsi-options target&lt;n&gt;-scsi-options scsi-reset-delay scsi-watchdog-tick scsi-initiator-id symFlags</pre> <p><code>target&lt;n&gt;-scsi-options</code> overrides the <code>scsi-options</code> property value for <code>target&lt;n&gt;</code>. <code>&lt;n&gt;</code> can vary from hex 0 to f. <code>symhisl</code> supports the following <code>scsi-options</code>: <code>SCSI_OPTIONS_DR</code>, <code>SCSI_OPTIONS_SYNC</code>, <code>SCSI_OPTIONS_FAST</code>, <code>SCSI_OPTIONS_ULTRA</code>, <code>SCSI_OPTIONS_ULTRA2</code>, <code>SCSI_OPTIONS_TAG</code>, and <code>SCSI_OPTIONS_WIDE</code>.</p> <p><code>SCSI_OPTIONS_PARITY</code> is supported for the <code>scsi-options</code> setting only and disables host adapter parity checking.</p> <p>After periodic interval <code>scsi-watchdog-tick</code> (seconds), the <code>symhisl</code> driver searches through all current and disconnected commands for timeouts.</p> <p><code>symFlags</code> is a driver-specific bit-mask you can use to enable or disable driver properties.</p> <table><tr><td>bit 0</td><td>When set, the driver will not reset the SCSI bus at initialization. Certain CD-ROM, tape, and other devices will not work properly when this bit is set. The default state for this bit is cleared.</td></tr><tr><td>bit 1</td><td>When set, the driver will not export the DMI ioctl interface. Set this bit only if you want to disable the ioctl interface for security reasons. The default state for this bit is cleared.</td></tr><tr><td>bit 2</td><td>When set, the driver disables 64-bit addressing capability. When clear, the driver enables 64-bit addressing capability. The default state for this bit is cleared.</td></tr></table> | bit 0 | When set, the driver will not reset the SCSI bus at initialization. Certain CD-ROM, tape, and other devices will not work properly when this bit is set. The default state for this bit is cleared. | bit 1 | When set, the driver will not export the DMI ioctl interface. Set this bit only if you want to disable the ioctl interface for security reasons. The default state for this bit is cleared. | bit 2 | When set, the driver disables 64-bit addressing capability. When clear, the driver enables 64-bit addressing capability. The default state for this bit is cleared. |
| bit 0                | When set, the driver will not reset the SCSI bus at initialization. Certain CD-ROM, tape, and other devices will not work properly when this bit is set. The default state for this bit is cleared.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |
| bit 1                | When set, the driver will not export the DMI ioctl interface. Set this bit only if you want to disable the ioctl interface for security reasons. The default state for this bit is cleared.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |
| bit 2                | When set, the driver disables 64-bit addressing capability. When clear, the driver enables 64-bit addressing capability. The default state for this bit is cleared.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                     |       |                                                                                                                                                                                             |       |                                                                                                                                                                     |

bit 3                    When set, the driver disables SCSI domain validation for all devices on any adapters controlled by the driver.

Refer to `scsi_hba_attach(9F)` for more information on driver configuration.

#### EXAMPLES

Edit the file `/kernel/drv/symhisl.conf` and add the following line:

```
scsi-options=0x78;
```

This disables tagged queuing, Fast, Ultra, and Ultra2 SCSI and wide mode for all `symhisl` instances.

The following example disables an option for one specific `symhisl` instance (refer to `driver.conf(4)` and `pci(4)` for more details):

```
name="symhisl" parent="/pci@1f,4000"
 unit-address="3"
 target1-scsi-options=0x58
 scsi-options=0x178 scsi-initiator-id=6;
```

Note that the initiator ID can only be changed for `symhisl` adapters that do not use the LSI Logic Boot ROM Configuration Utility. For adapters that can use the LSI Logic Boot ROM Configuration Utility, `scsi-initiator-id` has no effect.

The example above sets `scsi-options` for target 1 to 0x58 and all other targets on this SCSI bus to 0x178.

The physical path name of the parent can be determined using the `/devices` tree or following the link of the logical device name:

```
ls -l /dev/rdisk/c0t0d0s0
lrwxrwxrwx 1 root root 45 May 16 10:08 /dev/rdisk/c0t0d0s0 ->
 . . / . . /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

In this case, the parent is `/pci@1f,4000` and the unit-address is the number bound to the `scsi@3` node.

`scsi-options` specified per target ID have the highest precedence, followed by `scsi-options` per device type. Global `scsi-options` (for all `symhisl` instances) per bus have the lowest precedence.

The system must be rebooted for the specified `scsi-options` to take effect.

#### Driver Capabilities

The target driver sets capabilities in the `symhisl` driver to enable some driver features. The target driver can query and modify the following capabilities: `disconnect`, `synchronous`, `wide-xfer`, `tagged-qing`, and `auto-rqsense`. All other capabilities are query only.

By default, `tagged-qing` capabilities are disabled, while `disconnect`, `synchronous`, `wide-xfer`, `auto-rqsense`, and `untagged-qing` are enabled. These capabilities can only have binary values (0 or 1).

## symhis1(7D)

The target driver must enable `tagged-qing` explicitly. The `untagged-qing` capability is always enabled and its value cannot be modified.

If a conflict exists between the value of `scsi-options` and a capability, the value set in `scsi-options` prevails. Only whom `!= 0` is supported in the `scsi_ifsetcap(9F)` call. Refer to `scsi_ifsetcap(9F)` and `scsi_ifgetcap(9F)` for details.

**FILES** /kernel/drv/symhis1 ELF kernel module  
/kernel/drv/symhis1.conf Configuration file

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE                                                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Architecture   | Limited to PCI-based systems with LSI Logic SYM21002, SYM22910, SYM53C895A, SYM53C896, SYM53C1010-33, and SYM53C1010-66 SCSI I/O processors. |

**SEE ALSO** `prtconf(1M)`, `driver.conf(4)`, `pci(4)`, `attributes(5)`, `scsi_abort(9F)`, `scsi_hba_attach(9F)`, `scsi_ifgetcap(9F)`, `scsi_ifsetcap(9F)`, `scsi_reset(9F)`, `scsi_sync_pkt(9F)`, `scsi_transport(9F)`, `scsi_device(9S)`, `scsi_extended_sense(9S)`, `scsi_inquiry(9S)`, `scsi_pkt(9S)`

*Writing Device Drivers*

*ANSI Small Computer System Interface-2 (SCSI-2),*

*LSI Logic Corporation, SYM53C896 PCI-SCSI I/O Processor*

*LSI Logic Corporation, SYM53C895A PCI-SCSI I/O Processor*

*LSI Logic Corporation, SYM53C1010 PCI-SCSI I/O Processor*

**NOTES** The `symhis1` SYM53C895A and SYM53C896 (SYM21002 and SYM22910) hardware and software support Wide, Fast, SCSI Ultra, and Ultra2 synchronous speeds. SYM53C1010-33 and SYM53C1010-66 also support Ultra3 synchronous speeds. The maximum SCSI bandwidth for Ultra2 transfers is 80 Mbytes/sec and 160 Mbytes/sec for Ultra3.

- NAME** sysmsg – system message routing to console devices
- SYNOPSIS** /dev/sysmsg
- DESCRIPTION** The file /dev/sysmsg routes output to a variable set of console devices. Writes to /dev/sysmsg are always directed to the system console /dev/console, and are in addition directed to a set of auxiliary console devices managed by consadm(1M).
- Only root has permission to write to this device.
- ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Availability        | SUNWcsr         |
| Interface Stability | Stable          |

- SEE ALSO** consadm(1M), syslogd(1M), attributes(5), console(7D)

systrace(7D)

**NAME** systrace – DTrace system call tracing provider

**DESCRIPTION** The *systrace* driver implements the DTrace *syscall* dynamic tracing provider. The *syscall* provider performs dynamic instrumentation to offer probes that fire whenever a thread enters or returns from a kernel system call entry point.

The *systrace* driver is not a public interface and you access the instrumentation offered by this provider through DTrace. Refer to the *Solaris Dynamic Tracing Guide* for a description of the public documented interfaces available for the DTrace facility and the probes offered by the *syscall* provider.

**ATTRIBUTES** See *attributes(5)* for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Availability        | SUNWdtrp        |
| Interface Stability | Private         |

**SEE ALSO** *dtrace(1M)*, *attributes(5)*, *dtrace(7D)*  
*Solaris Dynamic Tracing Guide*

| <b>NAME</b>        | tavor – InfiniHost MT23108 InfiniBand (IB) Driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                |                 |              |                   |              |           |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-----------------|--------------|-------------------|--------------|-----------|
| <b>SYNOPSIS</b>    | PCI pci15b3,5a44@pci-slot, pci15b3,5a45@pci-slot                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                |                 |              |                   |              |           |
| <b>DESCRIPTION</b> | <p>The tavor driver is an IB Architecture-compliant implementation of an HCA, which operates on the Mellanox MT23108 InfiniBand ASSP. This ASSP supports the link and physical layers of the InfiniBand specification, while the ASSP and the driver support the transport layer.</p> <p>The tavor driver interfaces with the InfiniBand Transport Framework (IBTF) and provides an implementation of the Channel Interfaces that are defined by that framework. It also enables management applications and agents to access the IB fabric.</p>                                                                                                                                                         |                |                 |              |                   |              |           |
| <b>FILES</b>       | <p>/kernel/drv/tavor<br/>32-bit ELF kernel module (x86 platform only).</p> <p>/kernel/drv/amd64/tavor<br/>64-bit ELF kernel module (x86 platform only).</p> <p>/kernel/drv/sparcv9/tavor<br/>64-bit ELF Kernel Module (SPARC platform only).</p> <p>/kernel/drv/tavor.conf<br/>Driver configuration file.</p>                                                                                                                                                                                                                                                                                                                                                                                            |                |                 |              |                   |              |           |
| <b>ATTRIBUTES</b>  | See <code>attributes(5)</code> for descriptions of the following attributes:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                |                 |              |                   |              |           |
|                    | <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>PCI-based systems</td> </tr> <tr> <td>Availability</td> <td>SUNWtavor</td> </tr> </tbody> </table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Architecture | PCI-based systems | Availability | SUNWtavor |
| ATTRIBUTE TYPE     | ATTRIBUTE VALUE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                |                 |              |                   |              |           |
| Architecture       | PCI-based systems                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                |                 |              |                   |              |           |
| Availability       | SUNWtavor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                |                 |              |                   |              |           |
| <b>SEE ALSO</b>    | <p><code>driver.conf(4)</code>, <code>printers.conf(4)</code>, <code>attributes(5)</code></p> <p><i>Writing Device Drivers</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                |                 |              |                   |              |           |
| <b>DIAGNOSTICS</b> | <p>In addition to being logged, the following messages may appear on the system console.</p> <p>tavor : driver attached for maintenance mode only.<br/>There was a failure in the boot process of the tavor ASSP and the only function that can be performed is to re-flash firmware on the ASSP.</p> <p>driver failed to attach.<br/>The ASSP could not boot into either operational (HCA) mode or into maintenance mode. The device is inoperable.</p> <p>Unexpected port number in port state change event.<br/>A port state change event occurred, but the port number in the message does not exist on this HCA. This message also indicates the port number that was in the port state change.</p> |                |                 |              |                   |              |           |

tavor(7D)

Tavor driver successfully detached.

The driver has been removed from the system, and the HCA is no longer available for transfer operations.

tavor"n": port "m" up (link width 1x).

tavor"n": port "m" up (link width 4x).

tavor"n": port "m" up (link width unknown).

These messages occur when a port up asynchronous event occurs. "*n*" represents the instance of the Tavor device number, and "*m*" represents the port number on the Tavor device.

tavor"n": port "m" down.

A port down asynchronous event has occurred.

Tavor: <command name> command failed.

A internal firmware command failed to execute.



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | tcp, TCP – Internet Transmission Control Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>SYNOPSIS</b>    | <pre>#include &lt;sys/socket.h&gt;  #include &lt;netinet/in.h&gt;  s = socket(AF_INET, SOCK_STREAM, 0);  s = socket(AF_INET6, SOCK_STREAM, 0);  t = t_open("/dev/tcp", O_RDWR);  t = t_open("/dev/tcp6", O_RDWR);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>DESCRIPTION</b> | <p>TCP is the virtual circuit protocol of the Internet protocol family. It provides reliable, flow-controlled, in order, two-way transmission of data. It is a byte-stream protocol layered above the Internet Protocol (IP), or the Internet Protocol Version 6 (IPv6), the Internet protocol family's internetwork datagram delivery protocol.</p> <p>Programs can access TCP using the socket interface as a <code>SOCK_STREAM</code> socket type, or using the Transport Level Interface (TLI) where it supports the connection-oriented (<code>T_COTS_ORD</code>) service type.</p> <p>TCP uses IP's host-level addressing and adds its own per-host collection of "port addresses." The endpoints of a TCP connection are identified by the combination of an IP or IPv6 address and a TCP port number. Although other protocols, such as the User Datagram Protocol (UDP), may use the same host and port address format, the port space of these protocols is distinct. See <a href="#">inet(7P)</a> and <a href="#">inet6(7P)</a> for details on the common aspects of addressing in the Internet protocol family.</p> <p>Sockets utilizing TCP are either "active" or "passive." Active sockets initiate connections to passive sockets. Both types of sockets must have their local IP or IPv6 address and TCP port number bound with the <code>bind(3SOCKET)</code> system call after the socket is created. By default, TCP sockets are active. A passive socket is created by calling the <code>listen(3SOCKET)</code> system call after binding the socket with <code>bind()</code>. This establishes a queueing parameter for the passive socket. After this, connections to the passive socket can be received with the <code>accept(3SOCKET)</code> system call. Active sockets use the <code>connect(3SOCKET)</code> call after binding to initiate connections.</p> <p>By using the special value <code>INADDR_ANY</code> with IP, or the unspecified address (all zeroes) with IPv6, the local IP address can be left unspecified in the <code>bind()</code> call by either active or passive TCP sockets. This feature is usually used if the local address is either unknown or irrelevant. If left unspecified, the local IP or IPv6 address will be bound at connection time to the address of the network interface used to service the connection.</p> <p>Once a connection has been established, data can be exchanged using the <code>read(2)</code> and <code>write(2)</code> system calls.</p> <p>Under most circumstances, TCP sends data when it is presented. When outstanding data has not yet been acknowledged, TCP gathers small amounts of output to be sent in a single packet once an acknowledgement has been received. For a small number of clients, such as window systems that send a stream of mouse events which receive no</p> |

## tcp(7P)

replies, this packetization may cause significant delays. To circumvent this problem, TCP provides a socket-level boolean option, `TCP_NODELAY`. `TCP_NODELAY` is defined in `<netinet/tcp.h>`, and is set with `setsockopt(3SOCKET)` and tested with `getsockopt(3SOCKET)`. The option level for the `setsockopt()` call is the protocol number for TCP, available from `getprotobyname(3SOCKET)`.

Another socket level option, `SO_RCVBUF`, can be used to control the window that TCP advertises to the peer. IP level options may also be used with TCP. See [ip\(7P\)](#) and [ip6\(7p\)](#).

TCP provides an urgent data mechanism, which may be invoked using the out-of-band provisions of `send(3SOCKET)`. The caller may mark one byte as “urgent” with the `MSG_OOB` flag to `send(3SOCKET)`. This sets an “urgent pointer” pointing to this byte in the TCP stream. The receiver on the other side of the stream is notified of the urgent data by a `SIGURG` signal. The `SIOCATMARK ioctl(2)` request returns a value indicating whether the stream is at the urgent mark. Because the system never returns data across the urgent mark in a single `read(2)` call, it is possible to advance to the urgent data in a simple loop which reads data, testing the socket with the `SIOCATMARK ioctl()` request, until it reaches the mark.

Incoming connection requests that include an IP source route option are noted, and the reverse source route is used in responding.

A checksum over all data helps TCP implement reliability. Using a window-based flow control mechanism that makes use of positive acknowledgements, sequence numbers, and a retransmission strategy, TCP can usually recover when datagrams are damaged, delayed, duplicated or delivered out of order by the underlying communication medium.

If the local TCP receives no acknowledgements from its peer for a period of time, as would be the case if the remote machine crashed, the connection is closed and an error is returned to the user. If the remote machine reboots or otherwise loses state information about a TCP connection, the connection is aborted and an error is returned to the user.

SunOS supports TCP Extensions for High Performance (RFC 1323) which includes the window scale and time stamp options, and Protection Against Wrap Around Sequence Numbers (PAWS). SunOS also supports Selective Acknowledgment (SACK) capabilities (RFC 2018) and Explicit Congestion Notification (ECN) mechanism (RFC 3168).

Turn on the window scale option in one of the following ways:

- An application can set `SO_SNDBUF` or `SO_RCVBUF` size in the `setsockopt()` option to be larger than 64K. This must be done *before* the program calls `listen()` or `connect()`, because the window scale option is negotiated when the connection is established. Once the connection has been made, it is too late to increase the send or receive window beyond the default TCP limit of 64K.

- For all applications, use `ndd(1M)` to modify the configuration parameter `tcp_wscale_always`. If `tcp_wscale_always` is set to 1, the window scale option will always be set when connecting to a remote system. If `tcp_wscale_always` is 0, the window scale option will be set only if the user has requested a send or receive window larger than 64K. The default value of `tcp_wscale_always` is 0.
- Regardless of the value of `tcp_wscale_always`, the window scale option will always be included in a connect acknowledgement if the connecting system has used the option.

Turn on SACK capabilities in the following way:

- Use `ndd` to modify the configuration parameter `tcp_sack_permitted`. If `tcp_sack_permitted` is set to 0, TCP will not accept SACK or send out SACK information. If `tcp_sack_permitted` is set to 1, TCP will not initiate a connection with SACK permitted option in the SYN segment, but will respond with SACK permitted option in the SYN|ACK segment if an incoming connection request has the SACK permitted option. This means that TCP will only accept SACK information if the other side of the connection also accepts SACK information. If `tcp_sack_permitted` is set to 2, it will both initiate and accept connections with SACK information. The default for `tcp_sack_permitted` is 2 (active enabled).

Turn on TCP ECN mechanism in the following way:

- Use `ndd` to modify the configuration parameter `tcp_ecn_permitted`. If `tcp_ecn_permitted` is set to 0, TCP will not negotiate with a peer that supports ECN mechanism. If `tcp_ecn_permitted` is set to 1 when initiating a connection, TCP will not tell a peer that it supports ECN mechanism. However, it will tell a peer that it supports ECN mechanism when accepting a new incoming connection request if the peer indicates that it supports ECN mechanism in the SYN segment. If `tcp_ecn_permitted` is set to 2, in addition to negotiating with a peer on ECN mechanism when accepting connections, TCP will indicate in the outgoing SYN segment that it supports ECN mechanism when TCP makes active outgoing connections. The default for `tcp_ecn_permitted` is 1.

Turn on the time stamp option in the following way:

- Use `ndd` to modify the configuration parameter `tcp_tstamp_always`. If `tcp_tstamp_always` is 1, the time stamp option will always be set when connecting to a remote machine. If `tcp_tstamp_always` is 0, the timestamp option will not be set when connecting to a remote system. The default for `tcp_tstamp_always` is 0.
- Regardless of the value of `tcp_tstamp_always`, the time stamp option will always be included in a connect acknowledgement (and all succeeding packets) if the connecting system has used the time stamp option.

Use the following procedure to turn on the time stamp option only when the window scale option is in effect:

## tcp(7P)

- Use `ndd` to modify the configuration parameter `tcp_tstamp_if_wscale`. Setting `tcp_tstamp_if_wscale` to 1 will cause the time stamp option to be set when connecting to a remote system, if the window scale option has been set. If `tcp_tstamp_if_wscale` is 0, the time stamp option will not be set when connecting to a remote system. The default for `tcp_tstamp_if_wscale` is 1.

Protection Against Wrap Around Sequence Numbers (PAWS) is always used when the time stamp option is set.

SunOS also supports multiple methods of generating initial sequence numbers. One of these methods is the improved technique suggested in RFC 1948. We *HIGHLY* recommend that you set sequence number generation parameters to be as close to boot time as possible. This prevents sequence number problems on connections that use the same connection-ID as ones that used a different sequence number generation. The `svc:/network/initial:default` service configures the initial sequence number generation. The service reads the value contained in the configuration file `/etc/default/inetinit` to determine which method to use.

The `/etc/default/inetinit` file is an unstable interface, and may change in future releases.

TCP may be configured to report some information on connections that terminate by means of an RST packet. By default, no logging is done. If the `ndd(1M)` parameter `tcp_trace` is set to 1, then trace data is collected for all new connections established after that time.

The trace data consists of the TCP headers and IP source and destination addresses of the last few packets sent in each direction before RST occurred. Those packets are logged in a series of `strlog(9F)` calls. This trace facility has a very low overhead, and so is superior to such utilities as `snoop(1M)` for non-intrusive debugging for connections terminating by means of an RST.

**SEE ALSO** [svcs\(1\)](#), [ndd\(1M\)](#), [ioctl\(2\)](#), [read\(2\)](#), [svcadm\(1M\)](#), [write\(2\)](#), [accept\(3SOCKET\)](#), [bind\(3SOCKET\)](#), [connect\(3SOCKET\)](#), [getprotobyname\(3SOCKET\)](#), [getsockopt\(3SOCKET\)](#), [listen\(3SOCKET\)](#), [send\(3SOCKET\)](#), [smf\(5\)](#), [inet\(7P\)](#), [inet6\(7P\)](#), [ip\(7P\)](#), [ip6\(7P\)](#)

Ramakrishnan, K., Floyd, S., Black, D., RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, September 2001.

Mathias, M. and Hahdavi, J. Pittsburgh Supercomputing Center; Ford, S. Lawrence Berkeley National Laboratory; Romanow, A. Sun Microsystems, Inc. RFC 2018, *TCP Selective Acknowledgement Options*, October 1996.

Bellovin, S., RFC 1948, *Defending Against Sequence Number Attacks*, May 1996.

Jacobson, V., Braden, R., and Borman, D., RFC 1323, *TCP Extensions for High Performance*, May 1992.

Postel, Jon, *RFC 793, Transmission Control Protocol - DARPA Internet Program Protocol Specification*, Network Information Center, SRI International, Menlo Park, CA., September 1981.

**DIAGNOSTICS**

A socket operation may fail if:

|               |                                                                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| EISCONN       | A <code>connect()</code> operation was attempted on a socket on which a <code>connect()</code> operation had already been performed.                    |
| ETIMEDOUT     | A connection was dropped due to excessive retransmissions.                                                                                              |
| ECONNRESET    | The remote peer forced the connection to be closed (usually because the remote machine has lost state information about the connection due to a crash). |
| ECONNREFUSED  | The remote peer actively refused connection establishment (usually because no process is listening to the port).                                        |
| EADDRINUSE    | A <code>bind()</code> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.               |
| EADDRNOTAVAIL | A <code>bind()</code> operation was attempted on a socket with a network address for which no network interface exists.                                 |
| EACCES        | A <code>bind()</code> operation was attempted with a “reserved” port number and the effective user ID of the process was not the privileged user.       |
| ENOBUFS       | The system ran out of memory for internal data structures.                                                                                              |

**NOTES**

The `tcp` service is managed by the service management facility, `smf(5)`, under the service identifier:

```
svc:/network/initial:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using `svcadm(1M)`. The service’s status can be queried using the `svcs(1)` command.

## termio(7I)

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                             | termio – general terminal interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>SYNOPSIS</b>                         | <pre>#include &lt;termio.h&gt;  ioctl(int fildes, int request, struct termio *arg); ioctl(int fildes, int request, int arg);  #include &lt;termios.h&gt;  ioctl(int fildes, int request, struct termios *arg);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>DESCRIPTION</b>                      | <p>This release supports a general interface for asynchronous communications ports that is hardware-independent. The user interface to this functionality is using function calls (the preferred interface) described in <code>termios(3C)</code> or <code>ioctl</code> commands described in this section. This section also discusses the common features of the terminal subsystem which are relevant with both user interfaces.</p> <p>When a terminal file is opened, it normally causes the process to wait until a connection is established. In practice, user programs seldom open terminal files; they are opened by the system and become a user's standard input, output, and error files. The first terminal file opened by the session leader that is not already associated with a session becomes the controlling terminal for that session. The controlling terminal plays a special role in handling quit and interrupt signals, as discussed below. The controlling terminal is inherited by a child process during a <code>fork(2)</code>. A process can break this association by changing its session using <code>setsid()</code> (see <code>setsid(2)</code>).</p> <p>A terminal associated with one of these files ordinarily operates in full-duplex mode. Characters may be typed at any time, even while output is occurring, and are only lost when the character input buffers of the system become completely full, which is rare. For example, the number of characters in the line discipline buffer may exceed <code>{MAX_CANON}</code> and <code>IMAXBEL</code> (see below) is not set, or the user may accumulate <code>{MAX_INPUT}</code> number of input characters that have not yet been read by some program. When the input limit is reached, all the characters saved in the buffer up to that point are thrown away without notice.</p> |
| <b>Session Management (Job Control)</b> | <p>A control terminal will distinguish one of the process groups in the session associated with it to be the foreground process group. All other process groups in the session are designated as background process groups. This foreground process group plays a special role in handling signal-generating input characters, as discussed below. By default, when a controlling terminal is allocated, the controlling process's process group is assigned as foreground process group.</p> <p>Background process groups in the controlling process's session are subject to a job control line discipline when they attempt to access their controlling terminal. Process groups can be sent signals that will cause them to stop, unless they have made other arrangements. An exception is made for members of orphaned process groups.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

An orphaned process group is one where the process group (see `getpgid(2)`) has no members with a parent in a different process group but sharing the same controlling terminal. When a member of an orphaned process group attempts to access its controlling terminal, `EIO` is returned because there would be no way to restart the process if it were stopped on one of these signals.

If a member of a background process group attempts to read its controlling terminal, its process group will be sent a `SIGTTIN` signal, which will normally cause the members of that process group to stop. If, however, the process is ignoring or holding `SIGTTIN`, or is a member of an orphaned process group, the read will fail with `errno` set to `EIO`, and no signal is sent.

If a member of a background process group attempts to write its controlling terminal and the `TOSTOP` bit is set in the `c_lflag` field, its process group is sent a `SIGTTOU` signal, which will normally cause the members of that process group to stop. If, however, the process is ignoring or holding `SIGTTOU`, the write will succeed. If the process is not ignoring or holding `SIGTTOU` and is a member of an orphaned process group, the write will fail with `errno` set to `EIO`, and no signal will be sent.

If `TOSTOP` is set and a member of a background process group attempts to `ioctl` its controlling terminal, and that `ioctl` will modify terminal parameters (for example, `TCSETA`, `TCSETAW`, `TCSETAF`, or `TIOCSGRP`), its process group will be sent a `SIGTTOU` signal, which will normally cause the members of that process group to stop. If, however, the process is ignoring or holding `SIGTTOU`, the `ioctl` will succeed. If the process is not ignoring or holding `SIGTTOU` and is a member of an orphaned process group, the write will fail with `errno` set to `EIO`, and no signal will be sent.

### Canonical Mode Input Processing

Normally, terminal input is processed in units of lines. A line is delimited by a newline (ASCII LF) character, an end-of-file (ASCII EOT) character, or an end-of-line character. This means that a program attempting to read will block until an entire line has been typed. Also, no matter how many characters are requested in the read call, at most one line will be returned. It is not necessary, however, to read a whole line at once; any number of characters may be requested in a read, even one, without losing information.

During input, erase and kill processing is normally done. The ERASE character (by default, the character DEL) erases the last character typed. The WERASE character (the character `Control-w`) erases the last "word" typed in the current input line (but not any preceding spaces or tabs). A "word" is defined as a sequence of non-blank characters, with tabs counted as blanks. Neither ERASE nor WERASE will erase beyond the beginning of the line. The KILL character (by default, the character NAK) kills (deletes) the entire input line, and optionally outputs a newline character. All these characters operate on a key stroke basis, independent of any backspacing or tabbing that may have been done. The REPRINT character (the character `Control-r`) prints a newline followed by all characters that have not been read. Reprinting also occurs automatically if characters that would normally be erased from the screen are fouled by program output. The characters are reprinted as if they were being echoed; consequentially, if `ECHO` is not set, they are not printed.

termio(7I)

**Non-canonical  
Mode Input  
Processing**

The ERASE and KILL characters may be entered literally by preceding them with the escape character. In this case, the escape character is not read. The erase and kill characters may be changed.

In non-canonical mode input processing, input characters are not assembled into lines, and erase and kill processing does not occur. The MIN and TIME values are used to determine how to process the characters received.

MIN represents the minimum number of characters that should be received when the read is satisfied (that is, when the characters are returned to the user). TIME is a timer of 0.10-second granularity that is used to timeout bursty and short-term data transmissions. The four possible values for MIN and TIME and their interactions are described below.

Case A:  $\text{MIN} > 0, \text{TIME} > 0$

In this case, TIME serves as an intercharacter timer and is activated after the first character is received. Since it is an intercharacter timer, it is reset after a character is received. The interaction between MIN and TIME is as follows: as soon as one character is received, the intercharacter timer is started. If MIN characters are received before the intercharacter timer expires (note that the timer is reset upon receipt of each character), the read is satisfied. If the timer expires before MIN characters are received, the characters received to that point are returned to the user. Note that if TIME expires, at least one character will be returned because the timer would not have been enabled unless a character was received. In this case ( $\text{MIN} > 0, \text{TIME} > 0$ ), the read sleeps until the MIN and TIME mechanisms are activated by the receipt of the first character. If the number of characters read is less than the number of characters available, the timer is not reactivated and the subsequent read is satisfied immediately.

Case B:  $\text{MIN} > 0, \text{TIME} = 0$

In this case, since the value of TIME is zero, the timer plays no role and only MIN is significant. A pending read is not satisfied until MIN characters are received (the pending read sleeps until MIN characters



termio(7I)

are received). A program that uses this case to read record based terminal I/O may block indefinitely in the read operation.

Case C: MIN = 0, TIME > 0

In this case, since MIN = 0, TIME no longer represents an intercharacter timer: it now serves as a read timer that is activated as soon as a `read` is done. A read is satisfied as soon as a single character is received or the read timer expires. Note that, in this case, if the timer expires, no character is returned. If the timer does not expire, the only way the read can be satisfied is if a character is received. In this case, the read will not block indefinitely waiting for a character; if no character is received within TIME \*.10 seconds after the read is initiated, the read returns with zero characters.

Case D: MIN = 0, TIME = 0

In this case, return is immediate. The minimum of either the number of characters requested or the number of characters currently available is returned without waiting for more characters to be input.

### Comparing Different Cases of MIN, TIME Interaction

Some points to note about MIN and TIME :

- In the following explanations, note that the interactions of MIN and TIME are not symmetric. For example, when MIN > 0 and TIME = 0, TIME has no effect. However, in the opposite case, where MIN = 0 and TIME > 0, both MIN and TIME play a role in that MIN is satisfied with the receipt of a single character.
- Also note that in case A (MIN > 0, TIME > 0), TIME represents an intercharacter timer, whereas in case C ( MIN = 0, TIME > 0), TIME represents a read timer.

These two points highlight the dual purpose of the MIN/TIME feature. Cases A and B, where MIN > 0, exist to handle burst mode activity (for example, file transfer programs), where a program would like to process at least MIN characters at a time. In case A, the intercharacter timer is activated by a user as a safety measure; in case B, the timer is turned off.

Cases C and D exist to handle single character, timed transfers. These cases are readily adaptable to screen-based applications that need to know if a character is present in the input queue before refreshing the screen. In case C, the read is timed, whereas in case D, it is not.

Another important note is that MIN is always just a minimum. It does not denote a record length. For example, if a program does a read of 20 bytes, MIN is 10, and 25 characters are present, then 20 characters will be returned to the user.

termio(7I)

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Writing Characters</b> | When one or more characters are written, they are transmitted to the terminal as soon as previously written characters have finished typing. Input characters are echoed as they are typed if echoing has been enabled. If a process produces characters more rapidly than they can be typed, it will be suspended when its output queue exceeds some limit. When the queue is drained down to some threshold, the program is resumed.                                                                                                                                 |
| <b>Special Characters</b> | Certain characters have special functions on input. These functions and their default character values are summarized as follows:                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| INTR                      | (Control-c or ASCII ETX) generates a SIGINT signal. SIGINT is sent to all foreground processes associated with the controlling terminal. Normally, each such process is forced to terminate, but arrangements may be made either to ignore the signal or to receive a trap to an agreed upon location. (See <code>signal.h(3HEAD)</code> ).                                                                                                                                                                                                                            |
| QUIT                      | (Control-l or ASCII FS) generates a SIGQUIT signal. Its treatment is identical to the interrupt signal except that, unless a receiving process has made other arrangements, it will not only be terminated but a core image file (called <code>core</code> ) will be created in the current working directory.                                                                                                                                                                                                                                                         |
| ERASE                     | (DEL) erases the preceding character. It does not erase beyond the start of a line, as delimited by a NL, EOF, EOL, or EOL2 character.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| WERASE                    | (Control-w or ASCII ETX) erases the preceding "word". It does not erase beyond the start of a line, as delimited by a NL, EOF, EOL, or EOL2 character.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| KILL                      | (Control-u or ASCII NAK) deletes the entire line, as delimited by a NL, EOF, EOL, or EOL2 character.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| REPRINT                   | (Control-r or ASCII DC2) reprints all characters, preceded by a newline, that have not been read.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| EOF                       | (Control-d or ASCII EOT) may be used to generate an end-of-file from a terminal. When received, all the characters waiting to be read are immediately passed to the program, without waiting for a newline, and the EOF is discarded. Thus, if no characters are waiting (that is, the EOF occurred at the beginning of a line) zero characters are passed back, which is the standard end-of-file indication. Unless escaped, the EOF character is not echoed. Because EOT is the default EOF character, this prevents terminals that respond to EOT from hanging up. |
| NL                        | (ASCII LF) is the normal line delimiter. It cannot be changed or escaped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| EOL                       | (ASCII NULL) is an additional line delimiter, like NL. It is not normally used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| EOL2                      | is another additional line delimiter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

|         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SWTCH   | (Control-z or ASCII EM) Header file symbols related to this special character are present for compatibility purposes only and the kernel takes no special action on matching SWTCH (except to discard the character).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SUSP    | (Control-z or ASCII SUB) generates a SIGTSTP signal. SIGTSTP stops all processes in the foreground process group for that terminal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| DSUSP   | (Control-y or ASCII EM). It generates a SIGTSTP signal as SUSP does, but the signal is sent when a process in the foreground process group attempts to read the DSUSP character, rather than when it is typed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| STOP    | (Control-s or ASCII DC3) can be used to suspend output temporarily. It is useful with CRT terminals to prevent output from disappearing before it can be read. While output is suspended, STOP characters are ignored and not read.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| START   | (Control-q or ASCII DC1) is used to resume output. Output has been suspended by a STOP character. While output is not suspended, START characters are ignored and not read.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| DISCARD | (Control-o or ASCII SI) causes subsequent output to be discarded. Output is discarded until another DISCARD character is typed, more input arrives, or the condition is cleared by a program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| LNEXT   | (Control-v or ASCII SYN) causes the special meaning of the next character to be ignored. This works for all the special characters mentioned above. It allows characters to be input that would otherwise be interpreted by the system (for example KILL, QUIT). The character values for INTR, QUIT, ERASE, WERASE, KILL, REPRINT, EOF, EOL, EOL2, SWTCH, SUSP, DSUSP, STOP, START, DISCARD, and LNEXT may be changed to suit individual tastes. If the value of a special control character is _POSIX_VDISABLE (0), the function of that special control character is disabled. The ERASE, KILL, and EOF characters may be escaped by a preceding backslash (\) character, in which case no special function is done. Any of the special characters may be preceded by the LNEXT character, in which case no special function is done. |

### Modem Disconnect

When a modem disconnect is detected, a SIGHUP signal is sent to the terminal's controlling process. Unless other arrangements have been made, these signals cause the process to terminate. If SIGHUP is ignored or caught, any subsequent read returns with an end-of-file indication until the terminal is closed.

If the controlling process is not in the foreground process group of the terminal, a SIGTSTP is sent to the terminal's foreground process group. Unless other arrangements have been made, these signals cause the processes to stop.

## termio(7I)

Processes in background process groups that attempt to access the controlling terminal after modem disconnect while the terminal is still allocated to the session will receive appropriate SIGTTOU and SIGTTIN signals. Unless other arrangements have been made, this signal causes the processes to stop.

The controlling terminal will remain in this state until it is reinitialized with a successful open by the controlling process, or deallocated by the controlling process.

### Terminal Parameters

The parameters that control the behavior of devices and modules providing the `termios` interface are specified by the `termios` structure defined by `termios.h`. Several `ioctl(2)` system calls that fetch or change these parameters use this structure that contains the following members:

```
tcflag_t c_iflag; /* input modes */
tcflag_t c_oflag; /* output modes */
tcflag_t c_cflag; /* control modes */
tcflag_t c_lflag; /* local modes */
cc_t c_cc[NCCS]; /* control chars */
```

The special control characters are defined by the array `c_cc`. The symbolic name `NCCS` is the size of the Control-character array and is also defined by `<termios.h>`. The relative positions, subscript names, and typical default values for each function are as follows:

| Relative Position | Subscript Name | Typical Default Value |
|-------------------|----------------|-----------------------|
| 0                 | VINTR          | ETX                   |
| 1                 | VQUIT          | FS                    |
| 2                 | VERASE         | DEL                   |
| 3                 | VKILL          | NAK                   |
| 4                 | VEOF           | EOT                   |
| 5                 | VEOL           | NUL                   |
| 6                 | VEOL2          | NUL                   |
| 7                 | VWSTCH         | NUL                   |
| 8                 | VSTART         | NUL                   |
| 9                 | VSTOP          | DC3                   |
| 10                | VSUSP          | SUB                   |
| 11                | VDSUSP         | EM                    |
| 12                | VREPRINT       | DC2                   |
| 13                | VDISCARD       | SI                    |

| Relative Position | Subscript Name | Typical Default Value |
|-------------------|----------------|-----------------------|
| 14                | VWERASE        | ETB                   |
| 15                | VLNEXT         | SYN                   |
| 16-19             | Reserved       |                       |

**Input Modes**

The `c_iflag` field describes the basic terminal input control:

|         |                                         |
|---------|-----------------------------------------|
| IGNBRK  | Ignore break condition.                 |
| BRKINT  | Signal interrupt on break.              |
| IGNPAR  | Ignore characters with parity errors.   |
| PARMRK  | Mark parity errors.                     |
| INPCK   | Enable input parity check.              |
| ISTRIP  | Strip character.                        |
| INLCR   | Map NL to CR on input.                  |
| IGNCR   | Ignore CR.                              |
| ICRNL   | Map CR to NL on input.                  |
| IUCLC   | Map upper-case to lower-case on input.  |
| IXON    | Enable start/stop output control.       |
| IXANY   | Enable any character to restart output. |
| IXOFF   | Enable start/stop input control.        |
| IMAXBEL | Echo BEL on input line too long.        |

If `IGNBRK` is set, a break condition (a character framing error with data all zeros) detected on input is ignored, that is, not put on the input queue and therefore not read by any process. If `IGNBRK` is not set and `BRKINT` is set, the break condition shall flush the input and output queues and if the terminal is the controlling terminal of a foreground process group, the break condition generates a single `SIGINT` signal to that foreground process group. If neither `IGNBRK` nor `BRKINT` is set, a break condition is read as a single `'\0'` (ASCII NULL) character, or if `PARMRK` is set, as `'\377'`, `'\0'`, `c`, where `'\377'` is a single character with value 377 octal (0xff hex, 255 decimal), `'\0'` is a single character with value 0, and `c` is the errored character received.

If `IGNPAR` is set, a byte with framing or parity errors (other than break) is ignored.

## termio(7I)

If PARMRK is set, and IGNPAR is not set, a byte with a framing or parity error (other than break) is given to the application as the three-character sequence: `'\377', '\0', c`, where `'\377'` is a single character with value 377 octal (0xff hex, 255 decimal), `'\0'` is a single character with value 0, and `c` is the errored character received. To avoid ambiguity in this case, if ISTRIP is not set, a valid character of `'\377'` is given to the application as `'\377.'` If neither IGNPAR nor PARMRK is set, a framing or parity error (other than break) is given to the application as a single `'\0'` (ASCII NULL) character.

If INPCK is set, input parity checking is enabled. If INPCK is not set, input parity checking is disabled. This allows output parity generation without input parity errors. Note that whether input parity checking is enabled or disabled is independent of whether parity detection is enabled or disabled. If parity detection is enabled but input parity checking is disabled, the hardware to which the terminal is connected will recognize the parity bit, but the terminal special file will not check whether this is set correctly or not.

If ISTRIP is set, valid input characters are first stripped to seven bits, otherwise all eight bits are processed.

If INLCR is set, a received NL character is translated into a CR character. If IGNCR is set, a received CR character is ignored (not read). Otherwise, if ICRNL is set, a received CR character is translated into a NL character.

If IUCLC is set, a received upper case, alphabetic character is translated into the corresponding lower case character.

If IXON is set, start/stop output control is enabled. A received STOP character suspends output and a received START character restarts output. The STOP and START characters will not be read, but will merely perform flow control functions. If IXANY is set, any input character restarts output that has been suspended.

If IXOFF is set, the system transmits a STOP character when the input queue is nearly full, and a START character when enough input has been read so that the input queue is nearly empty again.

If IMAXBEL is set, the ASCII BEL character is echoed if the input stream overflows. Further input is not stored, but any input already present in the input stream is not disturbed. If IMAXBEL is not set, no BEL character is echoed, and all input present in the input queue is discarded if the input stream overflows.

### Output Modes

The `c_oflag` field specifies the system treatment of output:

|       |                                    |
|-------|------------------------------------|
| OPOST | Post-process output.               |
| OLCUC | Map lower case to upper on output. |
| ONLCR | Map NL to CR-NL on output.         |
| OCRNL | Map CR to NL on output.            |
| ONOCR | No CR output at column 0.          |

|        |                                                |
|--------|------------------------------------------------|
| ONLRET | NL performs CR function.                       |
| OFILL  | Use fill characters for delay.                 |
| OFDEL  | Fill is DEL, else NULL.                        |
| NLDLY  | Select newline delays:                         |
|        | NL0                                            |
|        | NL1                                            |
| CRDLY  | Select carriage-return delays:                 |
|        | CR0                                            |
|        | CR1                                            |
|        | CR2                                            |
|        | CR3                                            |
| TABDLY | Select horizontal tab delays or tab expansion: |
|        | TAB0                                           |
|        | TAB1                                           |
|        | TAB2                                           |
|        | TAB3      Expand tabs to spaces                |
|        | XTABS     Expand tabs to spaces                |
| BSDLY  | Select backspace delays:                       |
|        | BS0                                            |
|        | BS1                                            |
| VTDLY  | Select vertical tab delays:                    |
|        | VT0                                            |
|        | VT1                                            |
| FFDLY  | Select form feed delays:                       |
|        | FF0                                            |
|        | FF1                                            |

If OPOST is set, output characters are post-processed as indicated by the remaining flags; otherwise, characters are transmitted without change.

If OLCUC is set, a lower case alphabetic character is transmitted as the corresponding upper case character. This function is often used in conjunction with IUCLC.

termio(7I)

If `ONLCR` is set, the `NL` character is transmitted as the `CR-NL` character pair. If `OCRNL` is set, the `CR` character is transmitted as the `NL` character. If `ONOCR` is set, no `CR` character is transmitted when at column 0 (first position). If `ONRET` is set, the `NL` character is assumed to do the carriage-return function; the column pointer is set to 0 and the delays specified for `CR` are used. Otherwise, the `NL` character is assumed to do just the line-feed function; the column pointer remains unchanged. The column pointer is also set to 0 if the `CR` character is actually transmitted.

The delay bits specify how long transmission stops to allow for mechanical or other movement when certain characters are sent to the terminal. In all cases, a value of 0 indicates no delay. If `OFILL` is set, fill characters are transmitted for delay instead of a timed delay. This is useful for high baud rate terminals that need only a minimal delay. If `OFDEL` is set, the fill character is `DEL`; otherwise it is `NULL`.

If a form-feed or vertical-tab delay is specified, it lasts for about 2 seconds.

Newline delay lasts about 0.10 seconds. If `ONLRET` is set, the carriage-return delays are used instead of the newline delays. If `OFILL` is set, two fill characters are transmitted.

Carriage-return delay type 1 is dependent on the current column position, type 2 is about 0.10 seconds, and type 3 is about 0.15 seconds. If `OFILL` is set, delay type 1 transmits two fill characters, and type 2 transmits four fill characters.

Horizontal-tab delay type 1 is dependent on the current column position. Type 2 is about 0.10 seconds. Type 3 specifies that tabs are to be expanded into spaces. If `OFILL` is set, two fill characters are transmitted for any delay.

Backspace delay lasts about 0.05 seconds. If `OFILL` is set, one fill character is transmitted.

The actual delays depend on line speed and system load.

#### Control Modes

The `c_cflag` field describes the hardware control of the terminal:

|                    |            |
|--------------------|------------|
| <code>CBAUD</code> | Baud rate: |
| <code>B0</code>    | Hang up    |
| <code>B50</code>   | 50 baud    |
| <code>B75</code>   | 75 baud    |
| <code>B110</code>  | 110 baud   |
| <code>B134</code>  | 134 baud   |
| <code>B150</code>  | 150 baud   |
| <code>B200</code>  | 200 baud   |
| <code>B300</code>  | 300 baud   |
| <code>B600</code>  | 600 baud   |



|          |                                                |
|----------|------------------------------------------------|
| B1200    | 1200 baud                                      |
| B1800    | 1800 baud                                      |
| B2400    | 2400 baud                                      |
| B4800    | 4800 baud                                      |
| B9600    | 9600 baud                                      |
| B19200   | 19200 baud                                     |
| EXTA     | External A                                     |
| B38400   | 38400 baud                                     |
| EXTB     | External B                                     |
| B57600   | 57600 baud                                     |
| B76800   | 76800 baud                                     |
| B115200  | 115200 baud                                    |
| B153600  | 153600 baud                                    |
| B230400  | 230400 baud                                    |
| B307200  | 307200 baud                                    |
| B460800  | 460800 baud                                    |
| CSIZE    | Character size:                                |
| CS5      | 5 bits                                         |
| CS6      | 6 bits                                         |
| CS7      | 7 bits                                         |
| CS8      | 8 bits                                         |
| CSTOPB   | Send two stop bits, else one                   |
| CREAD    | Enable receiver                                |
| PARENB   | Parity enable                                  |
| PARODD   | Odd parity, else even                          |
| HUPCL    | Hang up on last close                          |
| CLOCAL   | Local line, else dial-up                       |
| CIBAUD   | Input baud rate, if different from output rate |
| PAREXT   | Extended parity for mark and space parity      |
| CRTSXOFF | Enable inbound hardware flow control           |
| CRTSCTS  | Enable outbound hardware flow control          |

## termio(7I)

CBAUDEXT        Bit to indicate output speed > B38400

CIBAUDEXT       Bit to indicate input speed > B38400

The CBAUD bits together with the CBAUDEXT bit specify the output baud rate. To retrieve the output speed from the `termios` structure pointed to by `termios_p` see the following code segment.

```
speed_t ospeed;
if (termios_p->c_cflag & CBAUDEXT)
 ospeed = (termios_p->c_cflag & CBAUD) + CBAUD + 1;
else
 ospeed = termios_p->c_cflag & CBAUD;
```

To store the output speed in the `termios` structure pointed to by `termios_p` see the following code segment.

```
speed_t ospeed;
if (ospeed > CBAUD) {
 termios_p->c_cflag |= CBAUDEXT;
 ospeed -= (CBAUD + 1);
} else
 termios_p->c_cflag &= ~CBAUDEXT;
termios_p->c_cflag =
 (termios_p->c_cflag & ~CBAUD) | (ospeed & CBAUD);
```

The zero baud rate, B0, is used to hang up the connection. If B0 is specified, the data-terminal-ready signal is not asserted. Normally, this disconnects the line.

If the CIBAUDEXT or CIBAUD bits are not zero, they specify the input baud rate, with the CBAUDEXT and CBAUD bits specifying the output baud rate; otherwise, the output and input baud rates are both specified by the CBAUDEXT and CBAUD bits. The values for the CIBAUD bits are the same as the values for the CBAUD bits, shifted left `IBSHIFT` bits. For any particular hardware, impossible speed changes are ignored. To retrieve the input speed in the `termios` structure pointed to by `termios_p` see the following code segment.

```
speed_t ispeed;
if (termios_p->c_cflag & CIBAUDEXT)
 ispeed = ((termios_p->c_cflag & CIBAUD) >> IBSHIFT)
 + (CIBAUD >> IBSHIFT) + 1;
else
 ispeed = (termios_p->c_cflag & CIBAUD) >> IBSHIFT;
```

To store the input speed in the `termios` structure pointed to by `termios_p` see the following code segment.

```
speed_t ispeed;
if (ispeed == 0) {
 ispeed = termios_p->c_cflag & CBAUD;
if (termios_p->c_cflag & CBAUDEXT)
 ispeed += (CBAUD + 1);
```

```

}
if ((ispeed << IBSHIFT) > CIBAUD) {
 termios_p->c_cflag |= CIBAUDEXT;
 ispeed -= ((CIBAUD >> IBSHIFT) + 1);
} else
 termios_p->c_cflag &= ~CIBAUDEXT;
 termios_p->c_cflag =
 (termios_p->c_cflag & ~CIBAUD) |
 ((ispeed << IBSHIFT) & CIBAUD);

```

The CSIZE bits specify the character size in bits for both transmission and reception. This size does not include the parity bit, if any. If CSTOPB is set, two stop bits are used; otherwise, one stop bit is used. For example, at 110 baud, two stops bits are required.

If PARENB is set, parity generation and detection is enabled, and a parity bit is added to each character. If parity is enabled, the PARODD flag specifies odd parity if set; otherwise, even parity is used.

If CREAD is set, the receiver is enabled. Otherwise, no characters are received.

If HUPCL is set, the line is disconnected when the last process with the line open closes it or terminates. That is, the data-terminal-ready signal is not asserted.

If CLOCAL is set, the line is assumed to be a local, direct connection with no modem control; otherwise, modem control is assumed.

If CRTSXOFF is set, inbound hardware flow control is enabled.

If CRTSCTS is set, outbound hardware flow control is enabled.

The four possible combinations for the state of CRTSCTS and CRTSXOFF bits and their interactions are described below.

- |         |                                                                                                                                                                                                                                              |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Case A: | CRTSCTS off, CRTSXOFF off. In this case the hardware flow control is disabled.                                                                                                                                                               |
| Case B: | CRTSCTS on, CRTSXOFF off. In this case only outbound hardware flow control is enabled. The state of CTS signal is used to do outbound flow control. It is expected that output will be suspended if CTS is low and resumed when CTS is high. |
| Case C: | CRTSCTS off, CRTSXOFF on. In this case only inbound hardware flow control is enabled. The state of RTS signal is used to do inbound flow control. It is expected that input will be suspended if RTS is low and resumed when RTS is high.    |
| Case D: | CRTSCTS on, CRTSXOFF on. In this case both inbound and outbound hardware flow control are enabled. Uses the state of CTS signal to do outbound flow control and RTS signal to do inbound flow control.                                       |

## termio(7I)

### Local Modes

The `c_lflag` field of the argument structure is used by the line discipline to control terminal functions. The basic line discipline provides the following:

|         |                                                        |
|---------|--------------------------------------------------------|
| ISIG    | Enable signals.                                        |
| ICANON  | Canonical input (erase and kill processing).           |
| XCASE   | Canonical upper/lower presentation.                    |
| ECHO    | Enable echo.                                           |
| ECHOE   | Echo erase character as BS-SP-BS &.                    |
| ECHOK   | Echo NL after kill character.                          |
| ECHONL  | Echo NL .                                              |
| NOFLSH  | Disable flush after interrupt or quit.                 |
| TOSTOP  | Send SIGTTOU for background output.                    |
| ECHOCTL | Echo control characters as <i>char</i> , delete as ^?. |
| ECHOPRT | Echo erase character as character erased.              |
| ECHOKE  | BS-SP-BS erase entire line on line kill.               |
| FLUSHO  | Output is being flushed.                               |
| PENDIN  | Retype pending input at next read or input character.  |
| IEXTEN  | Enable extended (implementation-defined) functions.    |

If `ISIG` is set, each input character is checked against the special control characters `INTR`, `QUIT`, `SWTCH`, `SUSP`, `STATUS`, and `DSUSP`. If an input character matches one of these control characters, the function associated with that character is performed. (Note: If `SWTCH` is set and the character matches, the character is simply discarded. No other action is taken.) If `ISIG` is not set, no checking is done. Thus, these special input functions are possible only if `ISIG` is set.

If `ICANON` is set, canonical processing is enabled. This enables the erase and kill edit functions, and the assembly of input characters into lines delimited by `NL-c`, `EOF`, `EOL`, and `EOL .` If `ICANON` is not set, read requests are satisfied directly from the input queue. A read is not satisfied until at least `MIN` characters have been received or the timeout value `TIME` has expired between characters. This allows fast bursts of input to be read efficiently while still allowing single character input. The time value represents tenths of seconds.

If `XCASE` is set and `ICANON` is set, an upper case letter is accepted on input if preceded by a backslash (`\`) character, and is output preceded by a backslash (`\`) character. In this mode, the following escape sequences are generated on output and accepted on input:

| FOR: | USE: |
|------|------|
| '    | \'   |
|      | \    |
| ≈    | \^   |
| {    | \{   |
| }    | \}   |
| \    | \\   |

For example, input A as \a, \n as \\n, and \N as \\N.

If ECHO is set, characters are echoed as received.

When ICANON is set, the following echo functions are possible.

- If ECHO and ECHOE are set, and ECHOPRT is not set, the ERASE and WERASE characters are echoed as one or more ASCII BS SP BS, which clears the last character(s) from a CRT screen.
- If ECHO, ECHOPRT, and IEXTEN are set, the first ERASE and WERASE character in a sequence echoes as a backslash (\), followed by the characters being erased. Subsequent ERASE and WERASE characters echo the characters being erased, in reverse order. The next non-erase character causes a '/' (slash) to be typed before it is echoed. ECHOPRT should be used for hard copy terminals.
- If ECHOKE and IEXTEN are set, the kill character is echoed by erasing each character on the line from the screen (using the mechanism selected by ECHOE and ECHOPra).
- If ECHOK is set, and ECHOKE is not set, the NL character is echoed after the kill character to emphasize that the line is deleted. Note that a '\' (escape) character or an LNEXT character preceding the erase or kill character removes any special function.
- If ECHONL is set, the NL character is echoed even if ECHO is not set. This is useful for terminals set to local echo (so called half-duplex).

If ECHOCTL and IEXTEN are set, all control characters (characters with codes between 0 and 37 octal) other than ASCII TAB, ASCII NL, the START character, and the STOP character, ASCII CR, and ASCII BS are echoed as ^ X, where X is the character given by adding 100 octal to the code of the control character (so that the character with octal code 1 is echoed as ^ A) , and the ASCII DEL character, with code 177 octal, is echoed as ^ ?.

If NOFLSH is set, the normal flush of the input and output queues associated with the INTR, QUIT, and SUSP characters is not done. This bit should be set when restarting system calls that read from or write to a terminal (see `sigaction(2)`).

## termio(7I)

If TOSTOP and IEXTEN are set, the signal SIGTTOU is sent to a process that tries to write to its controlling terminal if it is not in the foreground process group for that terminal. This signal normally stops the process. Otherwise, the output generated by that process is output to the current output stream. Processes that are blocking or ignoring SIGTTOU signals are excepted and allowed to produce output, if any.

If FLUSHO and IEXTEN are set, data written to the terminal is discarded. This bit is set when the FLUSH character is typed. A program can cancel the effect of typing the FLUSH character by clearing FLUSHO.

If PENDIN and IEXTEN are set, any input that has not yet been read is reprinted when the next character arrives as input. PENDIN is then automatically cleared.

If IEXTEN is set, the following implementation-defined functions are enabled: special characters (WERASE, REPRINT, DISCARD, and LNEXT) and local flags (TOSTOP, ECHOCTL, ECHOPRT, ECHOKE, FLUSHO, and PENDIN).

### Minimum and Timeout

The MIN and TIME values were described previously, in the subsection, Non-canonical Mode Input Processing. The initial value of MIN is 1, and the initial value of TIME is 0.

### Terminal Size

The number of lines and columns on the terminal's display is specified in the winsize structure defined by `sys/termios.h` and includes the following members:

```
unsigned short ws_row; /* rows, in characters */
unsigned short ws_col; /* columns, in characters */
unsigned short ws_xpixel; /* horizontal size, in pixels */
unsigned short ws_ypixel; /* vertical size, in pixels */
```

### Termio Structure

The SunOS/SVR4 termio structure is used by some ioctls; it is defined by `sys/termio.h` and includes the following members:

```
unsigned short c_iflag; /* input modes */
unsigned short c_oflag; /* output modes */
unsigned short c_cflag; /* control modes */
unsigned short c_lflag; /* local modes */
char c_line; /* line discipline */
unsigned char c_cc[NCC]; /* control chars */
```

The special control characters are defined by the array `c_cc`. The symbolic name NCC is the size of the Control-character array and is also defined by `termio.h`. The relative positions, subscript names, and typical default values for each function are as follows:

| Relative Positions | Subscript Names | Typical Default Values |
|--------------------|-----------------|------------------------|
| 0                  | VINTR           | EXT                    |
| 1                  | VQUIT           | FS                     |

| Relative Positions | Subscript Names | Typical Default Values |
|--------------------|-----------------|------------------------|
| 2                  | VERASE          | DEL                    |
| 3                  | VKILL           | NAK                    |
| 4                  | VEOF            | EOT                    |
| 5                  | VEOL            | NUL                    |
| 6                  | VEOL2           | NUL                    |
| 7                  | Reserved        |                        |

The MIN value is stored in the `VMIN` element of the `c_cc` array; the `TIME` value is stored in the `VTIME` element of the `c_cc` array. The `VMIN` element is the same element as the `VEOF` element; the `VTIME` element is the same element as the `VEOL` element.

The calls that use the `termio` structure only affect the flags and control characters that can be stored in the `termio` structure; all other flags and control characters are unaffected.

### Modem Lines

On special files representing serial ports, modem control lines can be read. Control lines (if the underlying hardware supports it) may also be changed. Status lines are read-only. The following modem control and status lines may be supported by a device; they are defined by `sys/termios.h`:

|                        |                     |
|------------------------|---------------------|
| <code>TIOCM_LE</code>  | line enable         |
| <code>TIOCM_DTR</code> | data terminal ready |
| <code>TIOCM_RTS</code> | request to send     |
| <code>TIOCM_ST</code>  | secondary transmit  |
| <code>TIOCM_SR</code>  | secondary receive   |
| <code>TIOCM_CTS</code> | clear to send       |
| <code>TIOCM_CAR</code> | carrier detect      |
| <code>TIOCM_RNG</code> | ring                |
| <code>TIOCM_DSR</code> | data set ready      |

`TIOCM_CD` is a synonym for `TIOCM_CAR`, and `TIOCM_RI` is a synonym for `TIOCM_RNG`. Not all of these are necessarily supported by any particular device; check the manual page for the device in question.

The software carrier mode can be enabled or disabled using the `TIOCSSOFTCAR` `ioctl`. If the software carrier flag for a line is off, the line pays attention to the hardware carrier detect (DCD) signal. The `tty` device associated with the line cannot be opened until DCD is asserted. If the software carrier flag is on, the line behaves as if DCD is always asserted.

## termio(7I)

The software carrier flag is usually turned on for locally connected terminals or other devices, and is off for lines with modems.

To be able to issue the `TIOCGSOFTCAR` and `TIOCSSOFTCAR` `ioctl` calls, the `tty` line should be opened with `O_NDELAY` so that the `open(2)` will not wait for the carrier.

### Default Values

The initial `termios` values upon driver open is configurable. This is accomplished by setting the “`ttymodes`” property in the file `/kernel/drv/options.conf`. Since this property is assigned during system initialization, any change to the “`ttymodes`” property will not take effect until the next reboot. The string value assigned to this property should be in the same format as the output of the `stty(1)` command with the `-g` option.

If this property is undefined, the following `termios` modes are in effect. The initial input control value is `BRKINT`, `ICRNL`, `IXON`, `IMAXBEL`. The initial output control value is `OPOST`, `ONLCR`, `TAB3`. The initial hardware control value is `B9600`, `CS8`, `CREAD`. The initial line-discipline control value is `ISIG`, `ICANON`, `IEXTEN`, `ECHO`, `ECHOK`, `ECHOE`, `ECHOKE`, `ECHOCTL`.

### IOCTLS

The `ioctl`s supported by devices and `STREAMS` modules providing the `termios(3C)` interface are listed below. Some calls may not be supported by all devices or modules. The functionality provided by these calls is also available through the preferred function call interface specified on `termios`.

|                      |                                                                                                                                                                                                                                                                                                            |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>TCGETS</code>  | The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are fetched and stored into that structure.                                                                                                                                                                 |
| <code>TCSETS</code>  | The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are set from the values stored in that structure. The change is immediate.                                                                                                                                  |
| <code>TCSETSW</code> | The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that affect output.     |
| <code>TCSETSF</code> | The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs. |
| <code>TCGETA</code>  | The argument is a pointer to a <code>termio</code> structure. The current terminal parameters are fetched, and those parameters that can be stored in a <code>termio</code> structure are stored into that structure.                                                                                      |
| <code>TCSETA</code>  | The argument is a pointer to a <code>termio</code> structure. Those terminal parameters that can be stored in a <code>termio</code> structure are set from the values stored in that structure. The change is immediate.                                                                                   |



|            |                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TCSETAW    | The argument is a pointer to a <code>termio</code> structure. Those terminal parameters that can be stored in a <code>termio</code> structure are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that affect output.                                             |
| TCSETAF    | The argument is a pointer to a <code>termio</code> structure. Those terminal parameters that can be stored in a <code>termio</code> structure are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs.                                         |
| TCSBRK     | The argument is an <code>int</code> value. Wait for the output to drain. If the argument is 0, then send a break (zero valued bits for 0.25 seconds).                                                                                                                                                                                                                                             |
| TCXONC     | Start/stop control. The argument is an <code>int</code> value. If the argument is 0, suspend output; if 1, restart suspended output; if 2, suspend input; if 3, restart suspended input.                                                                                                                                                                                                          |
| TCFLSH     | The argument is an <code>int</code> value. If the argument is 0, flush the input queue; if 1, flush the output queue; if 2, flush both the input and output queues.                                                                                                                                                                                                                               |
| TIOCGPRP   | The argument is a pointer to a <code>pid_t</code> . Set the value of that <code>pid_t</code> to the process group ID of the foreground process group associated with the terminal. See <code>termios(3C)</code> for a description of <code>TCGETPGRP</code> .                                                                                                                                     |
| TIOCSPGRP  | The argument is a pointer to a <code>pid_t</code> . Associate the process group whose process group ID is specified by the value of that <code>pid_t</code> with the terminal. The new process group value must be in the range of valid process group ID values. Otherwise, the error <code>EPERM</code> is returned. See <code>termios(3C)</code> for a description of <code>TCSETPGRP</code> . |
| TIOCGSID   | The argument is a pointer to a <code>pid_t</code> . The session ID of the terminal is fetched and stored in the <code>pid_t</code> .                                                                                                                                                                                                                                                              |
| TIOCGWINSZ | The argument is a pointer to a <code>winsize</code> structure. The terminal driver's notion of the terminal size is stored into that structure.                                                                                                                                                                                                                                                   |
| TIOCSWINSZ | The argument is a pointer to a <code>winsize</code> structure. The terminal driver's notion of the terminal size is set from the values specified in that structure. If the new sizes are different from the old sizes, a <code>SIGWINCH</code> signal is set to the process group of the terminal.                                                                                               |
| TIOCMBIS   | The argument is a pointer to an <code>int</code> whose value is a mask containing modem control lines to be turned on. The control lines whose bits are set in the argument are turned on; no other control lines are affected.                                                                                                                                                                   |

## termio(7I)

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TIOCMBIC     | The argument is a pointer to an <code>int</code> whose value is a mask containing modem control lines to be turned off. The control lines whose bits are set in the argument are turned off; no other control lines are affected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| TIOCMGET     | The argument is a pointer to an <code>int</code> . The current state of the modem status lines is fetched and stored in the <code>int</code> pointed to by the argument.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| TIOCMSET     | The argument is a pointer to an <code>int</code> containing a new set of modem control lines. The modem control lines are turned on or off, depending on whether the bit for that mode is set or clear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| TIOCSPPS     | The argument is a pointer to an <code>int</code> that determines whether pulse-per-second event handling is to be enabled (non-zero) or disabled (zero). If a one-pulse-per-second reference clock is attached to the serial line's data carrier detect input, the local system clock will be calibrated to it. A clock with a high error, that is, a deviation of more than 25 microseconds per tick, is ignored.                                                                                                                                                                                                                                                                                                                            |
| TIOCGPPS     | The argument is a pointer to an <code>int</code> , in which the state of the even handling is returned. The <code>int</code> is set to a non-zero value if pulse-per-second (PPS) handling has been enabled. Otherwise, it is set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| TIOCGSOFTCAR | The argument is a pointer to an <code>int</code> whose value is 1 or 0, depending on whether the software carrier detect is turned on or off.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| TIOCSSOFTCAR | The argument is a pointer to an <code>int</code> whose value is 1 or 0. The value of the integer should be 0 to turn off software carrier, or 1 to turn it on.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| TIOCGPPSEV   | The argument is a pointer to a <code>struct ppsclockev</code> . This structure contains the following members:<br><br><pre>struct timeval tv;<br/>uint32_t serial;</pre> “tv” is the system clock timestamp when the event (pulse on the DCD pin) occurred. “serial” is the ordinal of the event, which each consecutive event being assigned the next ordinal. The first event registered gets a “serial” value of 1. The TIOCGPPSEV returns the last event registered; multiple calls will persistently return the same event until a new one is registered. In addition to time stamping and saving the event, if it is of one-second period and of consistently high accuracy, the local system clock will automatically calibrate to it. |

**FILES** Files in or under `/dev`

**SEE ALSO** `stty(1)`, `fork(2)`, `getpgid(2)`, `getsid(2)`, `ioctl(2)`, `setsid(2)`, `sigaction(2)`, `signal(3C)`, `termios(3C)`, `signal.h(3HEAD)`, `streamio(7I)`

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                        | termiox – extended general terminal interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>DESCRIPTION</b>                 | <p>The extended general terminal interface supplements the <code>termio(7I)</code> general terminal interface by adding support for asynchronous hardware flow control, isochronous flow control and clock modes, and local implementations of additional asynchronous features. Some systems may not support all of these capabilities because of either hardware or software limitations. Other systems may not permit certain functions to be disabled. In these cases the appropriate bits will be ignored. See <code>&lt;sys/termiox.h&gt;</code> for your system to find out which capabilities are supported.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hardware Flow Control Modes</b> | <p>Hardware flow control supplements the <code>termio(7I)</code> <code>IXON</code>, <code>IXOFF</code>, and <code>IXANY</code> character flow control. Character flow control occurs when one device controls the data transfer of another device by the insertion of control characters in the data stream between devices. Hardware flow control occurs when one device controls the data transfer of another device using electrical control signals on wires (circuits) of the asynchronous interface. Isochronous hardware flow control occurs when one device controls the data transfer of another device by asserting or removing the transmit clock signals of that device. Character flow control and hardware flow control may be simultaneously set.</p> <p>In asynchronous, full duplex applications, the use of the Electronic Industries Association's EIA-232-D Request To Send (RTS) and Clear To Send (CTS) circuits is the preferred method of hardware flow control. An interface to other hardware flow control methods is included to provide a standard interface to these existing methods.</p> <p>The EIA-232-D standard specified only unidirectional hardware flow control - the Data Circuit-terminating Equipment or Data Communications Equipment (DCE) indicates to the Data Terminal Equipment (DTE) to stop transmitting data. The <code>termiox</code> interface allows both unidirectional and bidirectional hardware flow control; when bidirectional flow control is enabled, either the DCE or DTE can indicate to each other to stop transmitting data across the interface. Note: It is assumed that the asynchronous port is configured as a DTE. If the connected device is also a DTE and not a DCE, then DTE to DTE (for example, terminal or printer connected to computer) hardware flow control is possible by using a null modem to interconnect the appropriate data and control circuits.</p> |
| <b>Clock Modes</b>                 | <p>Isochronous communication is a variation of asynchronous communication whereby two communicating devices may provide transmit and/or receive clock signals to one another. Incoming clock signals can be taken from the baud rate generator on the local isochronous port controller, from CCITT V.24 circuit 114, Transmitter Signal Element Timing - DCE source (EIA-232-D pin 15), or from CCITT V.24 circuit 115, Receiver Signal Element Timing - DCE source (EIA-232-D pin 17). Outgoing clock signals can be sent on CCITT V.24 circuit 113, Transmitter Signal Element Timing - DTE source (EIA-232-D pin 24), on CCITT V.24 circuit 128, Receiver Signal Element Timing - DTE source (no EIA-232-D pin), or not sent at all.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## termiox(7I)

### Terminal Parameters

In terms of clock modes, traditional asynchronous communication is implemented simply by using the local baud rate generator as the incoming transmit and receive clock source and not outputting any clock signals.

The parameters that control the behavior of devices providing the `termiox` interface are specified by the `termiox` structure defined in the `<sys/termiox.h>` header. Several `ioctl(2)` system calls that fetch or change these parameters use this structure:

```
#define NFF 5
struct termiox {
 unsigned short x_hflag; /* hardware flow control modes */
 unsigned short x_cflag; /* clock modes */
 unsigned short x_rflag[NFF]; /* reserved modes */
 unsigned short x_sflag; /* spare local modes */
};
```

The `x_hflag` field describes hardware flow control modes:

|         |         |                                                   |
|---------|---------|---------------------------------------------------|
| RTSXOFF | 0000001 | Enable RTS hardware flow control on input.        |
| CTSXON  | 0000002 | Enable CTS hardware flow control on output.       |
| DTRXOFF | 0000004 | Enable DTR hardware flow control on input.        |
| CDXON   | 0000010 | Enable CD hardware flow control on output.        |
| ISXOFF  | 0000020 | Enable isochronous hardware flow control on input |

The EIA-232-D DTR and CD circuits are used to establish a connection between two systems. The RTS circuit is also used to establish a connection with a modem. Thus, both DTR and RTS are activated when an asynchronous port is opened. If DTR is used for hardware flow control, then RTS must be used for connectivity. If CD is used for hardware flow control, then CTS must be used for connectivity. Thus, RTS and DTR (or CTS and CD) cannot both be used for hardware flow control at the same time. Other mutual exclusions may apply, such as the simultaneous setting of the [termio\(7I\)](#) HUPCL and the `termiox` DTRXOFF bits, which use the DTE ready line for different functions.

Variations of different hardware flow control methods may be selected by setting the appropriate bits. For example, bidirectional RTS/CTS flow control is selected by setting both the RTSXOFF and CTSXON bits and bidirectional DTR/CTS flow control is selected by setting both the DTRXOFF and CTSXON. Modem control or unidirectional CTS hardware flow control is selected by setting only the CTSXON bit.

As previously mentioned, it is assumed that the local asynchronous port (for example, computer) is configured as a DTE. If the connected device (for example, printer) is also a DTE, it is assumed that the device is connected to the computer's asynchronous port using a null modem that swaps control circuits (typically RTS and CTS). The connected DTE drives RTS and the null modem swaps RTS and CTS so that the remote RTS is received as CTS by the local DTE. In the case that CTSXON is set for

hardware flow control, printer's lowering of its RTS would cause CTS seen by the computer to be lowered. Output to the printer is suspended until the printer's raising of its RTS, which would cause CTS seen by the computer to be raised.

If `RTSXOFF` is set, the Request To Send (RTS) circuit (line) will be raised, and if the asynchronous port needs to have its input stopped, it will lower the Request To Send (RTS) line. If the RTS line is lowered, it is assumed that the connected device will stop its output until RTS is raised.

If `CTSxon` is set, output will occur only if the Clear To Send (CTS) circuit (line) is raised by the connected device. If the CTS line is lowered by the connected device, output is suspended until CTS is raised.

If `DTRXOFF` is set, the DTE Ready (DTR) circuit (line) will be raised, and if the asynchronous port needs to have its input stopped, it will lower the DTE Ready (DTR) line. If the DTR line is lowered, it is assumed that the connected device will stop its output until DTR is raised.

If `CDxon` is set, output will occur only if the Received Line Signal Detector (CD) circuit (line) is raised by the connected device. If the CD line is lowered by the connected device, output is suspended until CD is raised.

If `ISXOFF` is set, and if the isochronous port needs to have its input stopped, it will stop the outgoing clock signal. It is assumed that the connected device is using this clock signal to create its output. Transmit and receive clock sources are programmed using the `x_cflag` fields. If the port is not programmed for external clock generation, `ISXOFF` is ignored. Output isochronous flow control is supported by appropriate clock source programming using the `x_cflag` field and enabled at the remote connected device.

The `x_cflag` field specifies the system treatment of clock modes.

|                     |         |                                                                                                                        |
|---------------------|---------|------------------------------------------------------------------------------------------------------------------------|
| <code>XMTCLK</code> | 0000007 | Transmit clock source:                                                                                                 |
| <code>XCIBRG</code> | 0000000 | Get transmit clock from internal baud rate generator.                                                                  |
| <code>XCTSET</code> | 0000001 | Get transmit clock from transmitter signal element timing (DCE source) lead, CCITT V.24 circuit 114, EIA-232-D pin 15. |
| <code>XCRSET</code> | 0000002 | Get transmit clock from receiver signal element timing (DCE source) lead, CCITT V.24 circuit 115, EIA-232-D pin 17.    |
| <code>RCVCLK</code> | 0000070 | Receive clock source:                                                                                                  |
| <code>RCIBRG</code> | 0000000 | Get receive clock from internal baud rate generator.                                                                   |

termiox(7I)

|           |         |                                                                                                                       |
|-----------|---------|-----------------------------------------------------------------------------------------------------------------------|
| RCTSET    | 0000010 | Get receive clock from transmitter signal element timing (DCE source) lead, CCITT V.24 circuit 114, EIA-232-D pin 15. |
| RCRSET    | 0000020 | Get receive clock from receiver signal element timing (DCE source) lead, CCITT V.24 circuit 115, EIA-232-D pin 17.    |
| TSETCLK   | 0000700 | Transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24, clock source:          |
| TSETCOFF  | 0000000 | TSET clock not provided.                                                                                              |
| TSETCRBRG | 0000100 | Output receive baud rate generator on circuit 113.                                                                    |
| TSETCTBRG | 0000200 | Output transmit baud rate generator on circuit 113                                                                    |
| TSETCTSET | 0000300 | Output transmitter signal element timing (DCE source) on circuit 113.                                                 |
| TSETCRSET | 0000400 | Output receiver signal element timing (DCE source) on circuit 113.                                                    |
| RSETCLK   | 0007000 | Receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin, clock source:             |
| RSETCOFF  | 0000000 | RSET clock not provided.                                                                                              |
| RSETCRBRG | 0001000 | Output receive baud rate generator on circuit 128.                                                                    |
| RSETCTBRG | 0002000 | Output transmit baud rate generator on circuit 128.                                                                   |
| RSETCTSET | 0003000 | Output transmitter signal element timing (DCE source) on circuit 128.                                                 |
| RSETCRSET | 0004000 | Output receiver signal element timing (DCE) on circuit 128.                                                           |

If the XMTCLK field has a value of XCIBRG the transmit clock is taken from the hardware internal baud rate generator, as in normal asynchronous transmission. If XMTCLK = XCTSET the transmit clock is taken from the Transmitter Signal Element Timing (DCE source) circuit. If XMTCLK = XCRSET the transmit clock is taken from the Receiver Signal Element Timing (DCE source) circuit.

If the RCVCLK field has a value of RCIBRG the receive clock is taken from the hardware Internal Baud Rate Generator, as in normal asynchronous transmission. If RCVCLK = RCTSET the receive clock is taken from the Transmitter Signal Element Timing (DCE source) circuit. If RCVCLK = RCRSET the receive clock is taken from the Receiver Signal Element Timing (DCE source) circuit.

If the TSETCLK field has a value of TSETCOFF the Transmitter Signal Element Timing (DTE source) circuit is not driven. If TSETCLK = TSETCRBRG the Transmitter Signal Element Timing (DTE source) circuit is driven by the Receive Baud Rate Generator. If TSETCLK = TSETCTBRG the Transmitter Signal Element Timing (DTE source) circuit is driven by the Transmit Baud Rate Generator. If TSETCLK = TSETCTSET the Transmitter Signal Element Timing (DTE source) circuit is driven by the Transmitter Signal Element Timing (DCE source). If TSETCLK = TSETCRBRG the Transmitter Signal Element Timing (DTE source) circuit is driven by the Receiver Signal Element Timing (DCE source).

If the RSETCLK field has a value of RSETCOFF the Receiver Signal Element Timing (DTE source) circuit is not driven. If RSETCLK = RSETCRBRG the Receiver Signal Element Timing (DTE source) circuit is driven by the Receive Baud Rate Generator. If RSETCLK = RSETCTBRG the Receiver Signal Element Timing (DTE source) circuit is driven by the Transmit Baud Rate Generator. If RSETCLK = RSETCTSET the Receiver Signal Element Timing (DTE source) circuit is driven by the Transmitter Signal Element Timing (DCE source). If RSETCLK = RSETCRBRG the Receiver Signal Element Timing (DTE source) circuit is driven by the Receiver Signal Element Timing (DCE source).

The `x_rflag` is reserved for future interface definitions and should not be used by any implementations. The `x_sflag` may be used by local implementations wishing to customize their terminal interface using the `termiox` `ioctl` system calls.

**IOCTLS** The `ioctl(2)` system calls have the form:

`ioctl (files, command, arg) struct termiox * arg;` The commands using this form are:

|         |                                                                                                                                                                                                                                                                                                             |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TCGETX  | The argument is a pointer to a <code>termiox</code> structure. The current terminal parameters are fetched and stored into that structure.                                                                                                                                                                  |
| TCSETX  | The argument is a pointer to a <code>termiox</code> structure. The current terminal parameters are set from the values stored in that structure. The change is immediate.                                                                                                                                   |
| TCSETXW | The argument is a pointer to a <code>termiox</code> structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that will affect output. |
| TCSETXF | The argument is a pointer to a <code>termiox</code> structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs.  |

**FILES** /dev/\*

**SEE ALSO** `stty(1)`, `ioctl(2)`, `termiox(7I)`

termiox(7I)

**NOTES** | The `termiox(7I)` system call is provided for compatibility with previous releases and its use is discouraged. Instead, the `termio(7I)` system call is recommended. See `termio(7I)` for usage information.



|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|---------------------------|-------------------------|--------------------------|-------------------------|-------|-------------------------------|---------------------|-------------------------|------------------------------------|--------------------------------|-----------------------------------|-------------------------------|---------------------|--------------------------------|---------------------------|------------------------------------|-------------------------------|--------------------------|------------------------------------|---------------------------------|-----------------------------------|--------------------------------|---------------------|---------------------------------|-------------------------------|
| <b>NAME</b>                        | ticlts, ticots, ticotsord – loopback transport providers                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <b>SYNOPSIS</b>                    | <pre>#include &lt;sys/ticlts.h&gt; #include &lt;sys/ticots.h&gt; #include &lt;sys/ticotsord.h&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <b>DESCRIPTION</b>                 | <p>The devices known as <code>ticlts</code>, <code>ticots</code>, and <code>ticotsord</code> are “loopback transport providers,” that is, stand-alone networks at the transport level. Loopback transport providers are transport providers in every sense except one: only one host (the local machine) is “connected to” a loopback network. Loopback transports present a TPI (STREAMS-level) interface to application processes and are intended to be accessed via the TLI (application-level) interface. They are implemented as clone devices and support address spaces consisting of “flex-addresses,” that is, arbitrary sequences of octets of length <math>&gt; 0</math>, represented by a <code>netbuf</code> structure.</p> <p><code>ticlts</code> is a datagram-mode transport provider. It offers (connectionless) service of type <code>T_CLTS</code>. Its default address size is <code>TCL_DEFAULTADDRSZ</code>. <code>ticlts</code> prints the following error messages (see <code>t_rcvuderr(3NSL)</code>):</p> <table border="0"> <tr> <td><code>TCL_BADADDR</code></td> <td>bad address specification</td> </tr> <tr> <td><code>TCL_BADOPT</code></td> <td>bad option specification</td> </tr> <tr> <td><code>TCL_NOPEER</code></td> <td>bound</td> </tr> <tr> <td><code>TCL_PEERBADSTATE</code></td> <td>peer in wrong state</td> </tr> </table> <p><code>ticots</code> is a virtual circuit-mode transport provider. It offers (connection-oriented) service of type <code>T_COTS</code>. Its default address size is <code>TCO_DEFAULTADDRSZ</code>. <code>ticots</code> prints the following disconnect messages (see <code>t_rcvdis(3NSL)</code>):</p> <table border="0"> <tr> <td><code>TCO_NOPEER</code></td> <td>no listener on destination address</td> </tr> <tr> <td><code>TCO_PEERNOROOMONQ</code></td> <td>peer has no room on connect queue</td> </tr> <tr> <td><code>TCO_PEERBADSTATE</code></td> <td>peer in wrong state</td> </tr> <tr> <td><code>TCO_PEERINITIATED</code></td> <td>peer-initiated disconnect</td> </tr> <tr> <td><code>TCO_PROVIDERINITIATED</code></td> <td>provider-initiated disconnect</td> </tr> </table> <p><code>ticotsord</code> is a virtual circuit-mode transport provider, offering service of type <code>T_COTS_ORD</code> (connection-oriented service with orderly release). Its default address size is <code>TCOO_DEFAULTADDRSZ</code>. <code>ticotsord</code> prints the following disconnect messages (see <code>t_rcvdis(3NSL)</code>):</p> <table border="0"> <tr> <td><code>TCOO_NOPEER</code></td> <td>no listener on destination address</td> </tr> <tr> <td><code>TCOO_PEERNOROOMONQ</code></td> <td>peer has no room on connect queue</td> </tr> <tr> <td><code>TCOO_PEERBADSTATE</code></td> <td>peer in wrong state</td> </tr> <tr> <td><code>TCOO_PEERINITIATED</code></td> <td>provider-initiated disconnect</td> </tr> </table> | <code>TCL_BADADDR</code> | bad address specification | <code>TCL_BADOPT</code> | bad option specification | <code>TCL_NOPEER</code> | bound | <code>TCL_PEERBADSTATE</code> | peer in wrong state | <code>TCO_NOPEER</code> | no listener on destination address | <code>TCO_PEERNOROOMONQ</code> | peer has no room on connect queue | <code>TCO_PEERBADSTATE</code> | peer in wrong state | <code>TCO_PEERINITIATED</code> | peer-initiated disconnect | <code>TCO_PROVIDERINITIATED</code> | provider-initiated disconnect | <code>TCOO_NOPEER</code> | no listener on destination address | <code>TCOO_PEERNOROOMONQ</code> | peer has no room on connect queue | <code>TCOO_PEERBADSTATE</code> | peer in wrong state | <code>TCOO_PEERINITIATED</code> | provider-initiated disconnect |
| <code>TCL_BADADDR</code>           | bad address specification                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCL_BADOPT</code>            | bad option specification                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCL_NOPEER</code>            | bound                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCL_PEERBADSTATE</code>      | peer in wrong state                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCO_NOPEER</code>            | no listener on destination address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCO_PEERNOROOMONQ</code>     | peer has no room on connect queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCO_PEERBADSTATE</code>      | peer in wrong state                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCO_PEERINITIATED</code>     | peer-initiated disconnect                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCO_PROVIDERINITIATED</code> | provider-initiated disconnect                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCOO_NOPEER</code>           | no listener on destination address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCOO_PEERNOROOMONQ</code>    | peer has no room on connect queue                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCOO_PEERBADSTATE</code>     | peer in wrong state                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |
| <code>TCOO_PEERINITIATED</code>    | provider-initiated disconnect                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                          |                           |                         |                          |                         |       |                               |                     |                         |                                    |                                |                                   |                               |                     |                                |                           |                                    |                               |                          |                                    |                                 |                                   |                                |                     |                                 |                               |

ticlts(7D)

TCOO\_PROVIDERINITIATED                      peer-initiated disconnect

**USAGE**      Loopback transports support a local IPC mechanism through the TLI interface. Applications implemented in a transport provider-independent manner on a client-server model using this IPC are transparently transportable to networked environments.

Transport provider-independent applications must not include the headers listed in the synopsis section above. In particular, the options are (like all transport provider options) provider dependent.

ticlts and ticots support the same service types (T\_CLTS and T\_COTS) supported by the OSI transport-level model.

ticotsord supports the same service type (T\_COTSORD) supported by the TCP/IP model.

**FILES**      /dev/ticlts

              /dev/ticots

              /dev/ticotsord

**SEE ALSO**    t\_rcvdis(3NSL), t\_rcvuderr(3NSL)

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |         |                                                                                                                                                                                                                                                                            |           |                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | timod – Transport Interface cooperating STREAMS module                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |         |                                                                                                                                                                                                                                                                            |           |                                                                                                                                                                                                                                                                                  |
| <b>SYNOPSIS</b>    | <pre>#include &lt;sys/stropts.h&gt;  ioctl(<i>fildev</i>, I_STR, &amp;<i>my_strioc1</i>);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |         |                                                                                                                                                                                                                                                                            |           |                                                                                                                                                                                                                                                                                  |
| <b>DESCRIPTION</b> | <p>timod is a STREAMS module for use with the Transport Interface (“TI”) functions of the Network Services library. The timod module converts a set of ioctl(2) calls into STREAMS messages that may be consumed by a transport protocol provider that supports the Transport Interface. This allows a user to initiate certain TI functions as atomic operations.</p> <p>The timod module must be pushed onto only a stream terminated by a transport protocol provider that supports the TI.</p> <p>All STREAMS messages, with the exception of the message types generated from the ioctl commands described below, will be transparently passed to the neighboring module or driver. The messages generated from the following ioctl commands are recognized and processed by the timod module. The format of the ioctl call is:</p> <pre>#include &lt;sys/stropts.h&gt; - - struct strioc1 <i>my_strioc1</i>; - - strioc1.ic_cmd = <i>cmd</i>; strioc1.ic_timeout = INFTIM; strioc1.ic_len = <i>size</i>; strioc1.ic_dp = (char *)<i>buf</i> ioctl(<i>fildev</i>, I_STR, &amp;<i>my_strioc1</i>);</pre> <p>On issuance, <i>size</i> is the size of the appropriate TI message to be sent to the transport provider and on return <i>size</i> is the size of the appropriate TI message from the transport provider in response to the issued TI message. <i>buf</i> is a pointer to a buffer large enough to hold the contents of the appropriate TI messages. The TI message types are defined in &lt;sys/tihdr.h&gt;. The possible values for the <i>cmd</i> field are:</p> <table border="0"> <tr> <td style="vertical-align: top; padding-right: 20px;">TI_BIND</td> <td>Bind an address to the underlying transport protocol provider. The message issued to the TI_BIND ioctl is equivalent to the TI message type T_BIND_REQ and the message returned by the successful completion of the ioctl is equivalent to the TI message type T_BIND_ACK.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">TI_UNBIND</td> <td>Unbind an address from the underlying transport protocol provider. The message issued to the TI_UNBIND ioctl is equivalent to the TI message type T_UNBIND_REQ and the message returned by the successful completion of the ioctl is equivalent to the TI message type T_OK_ACK.</td> </tr> </table> | TI_BIND | Bind an address to the underlying transport protocol provider. The message issued to the TI_BIND ioctl is equivalent to the TI message type T_BIND_REQ and the message returned by the successful completion of the ioctl is equivalent to the TI message type T_BIND_ACK. | TI_UNBIND | Unbind an address from the underlying transport protocol provider. The message issued to the TI_UNBIND ioctl is equivalent to the TI message type T_UNBIND_REQ and the message returned by the successful completion of the ioctl is equivalent to the TI message type T_OK_ACK. |
| TI_BIND            | Bind an address to the underlying transport protocol provider. The message issued to the TI_BIND ioctl is equivalent to the TI message type T_BIND_REQ and the message returned by the successful completion of the ioctl is equivalent to the TI message type T_BIND_ACK.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |         |                                                                                                                                                                                                                                                                            |           |                                                                                                                                                                                                                                                                                  |
| TI_UNBIND          | Unbind an address from the underlying transport protocol provider. The message issued to the TI_UNBIND ioctl is equivalent to the TI message type T_UNBIND_REQ and the message returned by the successful completion of the ioctl is equivalent to the TI message type T_OK_ACK.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |         |                                                                                                                                                                                                                                                                            |           |                                                                                                                                                                                                                                                                                  |

## timod(7M)

**TI\_GETINFO** Get the TI protocol specific information from the transport protocol provider. The message issued to the `TI_GETINFO` ioctl is equivalent to the TI message type `T_INFO_REQ` and the message returned by the successful completion of the `ioctl` is equivalent to the TI message type `T_INFO_ACK`.

**TI\_OPTMGMT** Get, set, or negotiate protocol specific options with the transport protocol provider. The message issued to the `TI_OPTMGMT` ioctl is equivalent to the TI message type `T_OPTMGMT_REQ` and the message returned by the successful completion of the `ioctl` is equivalent to the TI message type `T_OPTMGMT_ACK`.

**FILES**

|                                   |                                                                 |
|-----------------------------------|-----------------------------------------------------------------|
| <code>&lt;sys/timod.h&gt;</code>  | ioctl definitions                                               |
| <code>&lt;sys/tiuser.h&gt;</code> | TLI interface declaration and structure file                    |
| <code>&lt;sys/tihdr.h&gt;</code>  | TPI declarations and user-level code                            |
| <code>&lt;sys/errno.h&gt;</code>  | system error messages file. Please see <code>errno(3C)</code> . |

**SEE ALSO** `intro(3)`, `ioctl(2)`, `errno(3C)`, `tirdwr(7M)`

*STREAMS Programming Guide*

**DIAGNOSTICS** If the `ioctl` returns with a value greater than 0, the lower 8 bits of the return value will be one of the TI error codes as defined in `<sys/tiuser.h>`. If the TI error is of type `TSYSERR`, then the next 8 bits of the return value will contain an error as defined in `<sys/errno.h>` (see `intro(3)`).

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |      |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |      |                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | tirdwr – Transport Interface read/write interface STREAMS module                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |      |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |      |                                                                                                      |
| <b>SYNOPSIS</b>    | <pre>int ioctl( fd, I_PUSH, "tirdwr" );</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |      |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |      |                                                                                                      |
| <b>DESCRIPTION</b> | <p>tirdwr is a STREAMS module that provides an alternate interface to a transport provider which supports the Transport Interface (“TI”) functions of the Network Services library (see Section 3N). This alternate interface allows a user to communicate with the transport protocol provider using the <code>read(2)</code> and <code>write(2)</code> system calls. The <code>putmsg(2)</code> and <code>getmsg(2)</code> system calls may also be used. However, <code>putmsg</code> and <code>getmsg</code> can only transfer data messages between user and stream; control portions are disallowed.</p> <p>The tirdwr module must only be pushed (see <code>I_PUSH</code> in <a href="#">streamio(7I)</a>) onto a stream terminated by a transport protocol provider which supports the TI. After the tirdwr module has been pushed onto a stream, none of the TI functions can be used. Subsequent calls to TI functions cause an error on the stream. Once the error is detected, subsequent system calls on the stream return an error with <code>errno</code> set to <code>EPROTO</code>.</p> <p>The following are the actions taken by the tirdwr module when pushed on the stream, popped (see <code>I_POP</code> in <a href="#">streamio(7I)</a>) off the stream, or when data passes through it.</p> <table border="0"> <tr> <td style="vertical-align: top; padding-right: 20px;">push</td> <td>When the module is pushed onto a stream, it checks any existing data destined for the user to ensure that only regular data messages are present. It ignores any messages on the stream that relate to process management, such as messages that generate signals to the user processes associated with the stream. If any other messages are present, the <code>I_PUSH</code> will return an error with <code>errno</code> set to <code>EPROTO</code>.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">write</td> <td> <p>The module takes the following actions on data that originated from a <code>write</code> system call:</p> <ul style="list-style-type: none"> <li>■ All messages with the exception of messages that contain control portions (see the <code>putmsg</code> and <code>getmsg</code> system calls) are transparently passed onto the module’s downstream neighbor.</li> <li>■ Any zero length data messages are freed by the module and they will not be passed onto the module’s downstream neighbor.</li> <li>■ Any messages with control portions generate an error, and any further system calls associated with the stream fails with <code>errno</code> set to <code>EPROTO</code>.</li> </ul> </td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">read</td> <td>The module takes the following actions on data that originated from the transport protocol provider.</td> </tr> </table> | push | When the module is pushed onto a stream, it checks any existing data destined for the user to ensure that only regular data messages are present. It ignores any messages on the stream that relate to process management, such as messages that generate signals to the user processes associated with the stream. If any other messages are present, the <code>I_PUSH</code> will return an error with <code>errno</code> set to <code>EPROTO</code> . | write | <p>The module takes the following actions on data that originated from a <code>write</code> system call:</p> <ul style="list-style-type: none"> <li>■ All messages with the exception of messages that contain control portions (see the <code>putmsg</code> and <code>getmsg</code> system calls) are transparently passed onto the module’s downstream neighbor.</li> <li>■ Any zero length data messages are freed by the module and they will not be passed onto the module’s downstream neighbor.</li> <li>■ Any messages with control portions generate an error, and any further system calls associated with the stream fails with <code>errno</code> set to <code>EPROTO</code>.</li> </ul> | read | The module takes the following actions on data that originated from the transport protocol provider. |
| push               | When the module is pushed onto a stream, it checks any existing data destined for the user to ensure that only regular data messages are present. It ignores any messages on the stream that relate to process management, such as messages that generate signals to the user processes associated with the stream. If any other messages are present, the <code>I_PUSH</code> will return an error with <code>errno</code> set to <code>EPROTO</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |      |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |      |                                                                                                      |
| write              | <p>The module takes the following actions on data that originated from a <code>write</code> system call:</p> <ul style="list-style-type: none"> <li>■ All messages with the exception of messages that contain control portions (see the <code>putmsg</code> and <code>getmsg</code> system calls) are transparently passed onto the module’s downstream neighbor.</li> <li>■ Any zero length data messages are freed by the module and they will not be passed onto the module’s downstream neighbor.</li> <li>■ Any messages with control portions generate an error, and any further system calls associated with the stream fails with <code>errno</code> set to <code>EPROTO</code>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |      |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |      |                                                                                                      |
| read               | The module takes the following actions on data that originated from the transport protocol provider.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |      |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |      |                                                                                                      |

tirdwr(7M)

All messages with the exception of those that contain control portions (see the `putmsg` and `getmsg` system calls) are transparently passed onto the module's upstream neighbor. The action taken on messages with control portions will be as follows:

- Any data messages with control portions have the control portions removed from the message before to passing the message on to the upstream neighbor.
- Messages that represent an orderly release indication from the transport provider generate a zero length data message, indicating the end of file, which will be sent to the reader of the stream. The orderly release message itself is freed by the module.
- Messages that represent an abortive disconnect indication from the transport provider cause all further `write` and `putmsg` system calls to fail with `errno` set to `ENXIO`. All further `read` and `getmsg` system calls return zero length data (indicating end of file) once all previous data has been read.
- With the exception of the above rules, all other messages with control portions generate an error and all further system calls associated with the stream will fail with `errno` set to `EPROTO`.

Any zero length data messages are freed by the module and they are not passed onto the module's upstream neighbor.

pop

When the module is popped off the stream or the stream is closed, the module takes the following action:

- If an orderly release indication has been previously received, then an orderly release request will be sent to the remote side of the transport connection.

**SEE ALSO** [intro\(3\)](#), [getmsg\(2\)](#), [putmsg\(2\)](#), [read\(2\)](#), [write\(2\)](#), [intro\(3\)](#), [streamio\(7I\)](#), [timod\(7M\)](#)

*STREAMS Programming Guide*

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | tmpfs – memory based file system                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>SYNOPSIS</b>    | <pre>#include &lt;sys/mount.h&gt;  <b>mount</b> (<i>special, directory, MS_DATA, "tmpfs", NULL, 0</i>);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>DESCRIPTION</b> | <p>tmpfs is a memory based file system which uses kernel resources relating to the VM system and page cache as a file system. Once mounted, a tmpfs file system provides standard file operations and semantics. tmpfs is so named because files and directories are not preserved across reboot or unmounts, all files residing on a tmpfs file system that is unmounted will be lost.</p> <p>tmpfs file systems can be mounted with the command:</p> <pre>mount -F tmpfs swap <i>directory</i></pre> <p>Alternatively, to mount a tmpfs file system on /tmp at multi-user startup time (maximizing possible performance improvements), add the following line to /etc/vfstab:</p> <pre>swap -/tmp tmpfs - yes -</pre> <p>tmpfs is designed as a performance enhancement which is achieved by caching the writes to files residing on a tmpfs file system. Performance improvements are most noticeable when a large number of short lived files are written and accessed on a tmpfs file system. Large compilations with tmpfs mounted on /tmp are a good example of this.</p> <p>Users of tmpfs should be aware of some constraints involved in mounting a tmpfs file system. The resources used by tmpfs are the same as those used when commands are executed (for example, swap space allocation). This means that large sized tmpfs files can affect the amount of space left over for programs to execute. Likewise, programs requiring large amounts of memory use up the space available to tmpfs. Users running into this constraint (for example, running out of space on tmpfs) can allocate more swap space by using the swap(1M) command.</p> <p>Another constraint is that the number of files available in a tmpfs file system is calculated based on the physical memory of the machine and not the size of the swap device/partition. If you have too many files, tmpfs will print a warning message and you will be unable to create new files. You cannot increase this limit by adding swap space.</p> <p>Normal file system writes are scheduled to be written to a permanent storage medium along with all control information associated with the file (for example, modification time, file permissions). tmpfs control information resides only in memory and never needs to be written to permanent storage. File data remains in core until memory demands are sufficient to cause pages associated with tmpfs to be reused at which time they are copied out to swap.</p> |

## tmpfs(7FS)

An additional mount option can be specified to control the size of an individual tmpfs file system.

**SEE ALSO** `df(1M)`, `mount(1M)`, `mount_tmpfs(1M)`, `swap(1M)`, `mmap(2)`, `mount(2)`, `umount(2)`, `vfstab(4)`

*System Administration Guide: Basic Administration*

**DIAGNOSTICS** If tmpfs runs out of space, one of the following messages will display in the console.

*directory: File system full, swap space limit exceeded*

This message appears because a page could not be allocated while writing to a file.

This can occur if tmpfs is attempting to write more than it is allowed, or if currently executing programs are using a lot of memory. To make more space available, remove unnecessary files, exit from some programs, or allocate more swap space using `swap(1M)`.

*directory: File system full, memory allocation failed*

tmpfs ran out of physical memory while attempting to create a new file or directory. Remove unnecessary files or directories or install more physical memory.

**WARNINGS** Files and directories on a tmpfs file system are not preserved across reboots or unmounts. Command scripts or programs which count on this will not work as expected.

**NOTES** Compilers do not necessarily use `/tmp` to write intermediate files therefore missing some significant performance benefits. This can be remedied by setting the environment variable `TMPDIR` to `/tmp`. Compilers use the value in this environment variable as the name of the directory to store intermediate files.

swap to a tmpfs file is not supported.

`df(1M)` output is of limited accuracy since a tmpfs file system size is not static and the space available to tmpfs is dependent on the swap space demands of the entire system.



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | tokenmt – Single and Two Rate Three Conformance Level Meter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>DESCRIPTION</b> | <p>The tokenmt module can be configured as a Single or a Two Rate meter. Packets are deemed to belong to one of the three levels - Red, Yellow or Green - depending on the configured rate(s) and the burst sizes. When configured as a Single Rate meter, tokenmt can operate with just the Green and Red levels.</p> <p>Configuration parameters for tokenmt correspond to definitions in <i>RFC- 2697</i> and <i>RFC- 2698</i> as follows:</p> <p>Configuring tokenmt as a Single Rate meter (from <i>RFC- 2697</i>):</p> <p>committed_rate - CIR<br/> committed_burst - CBS<br/> peak_burst - EBS</p> <p>(thus 'peak_burst' for a single rate meter is actually the 'excess burst' in the RFC. However, throughout the text the parameter name "peak burst" is used.)</p> <p>Configuring tokenmt as a Two Rate meter (from <i>RFC- 2698</i>):</p> <p>committed_rate - CIR<br/> peak_rate - PIR<br/> committed_burst - CBS<br/> peak_burst - PBS</p> <p>The meter is implemented using token buckets C and P, which initially hold tokens equivalent to committed and peak burst sizes (bits) respectively. When a packet of size B bits arrive at time t, the following occurs:</p> <p>When operating as a Single Rate meter, the outcome (level) is decided as follows:</p> <ul style="list-style-type: none"> <li>- Update tokens in C and P <ul style="list-style-type: none"> <li>o Compute no. of tokens accumulated since the last time packet was seen at the committed rate as <math>T(t) = \text{committed rate} * (t - t')</math> (where t' is the time the last packet was seen)</li> <li>o Add T tokens to C up to a maximum of committed burst size. Add remaining tokens <math>((C+T) - \text{Committed Burst})</math>, if any, to P, to a maximum of peak burst size.</li> </ul> </li> <li>- Decide outcome <ul style="list-style-type: none"> <li>o If not color aware <ul style="list-style-type: none"> <li>o If <math>B \leq C</math>, outcome is GREEN and <math>C -= B</math>.</li> <li>o Else, if <math>B \leq P</math>, outcome is YELLOW and <math>P -= B</math>.</li> <li>o Else, outcome is Red.</li> </ul> </li> <li>o Else, <ul style="list-style-type: none"> <li>o obtain DSCP from packet</li> <li>o obtain color from color_map, color_map[DSCP]</li> <li>o if (color is GREEN) and <math>(B \leq C)</math>, outcome is GREEN and <math>C -= B</math>.</li> <li>o Else, if (color is GREEN or YELLOW) and <math>(B \leq P)</math>, outcome is YELLOW and <math>P -= B</math>.</li> </ul> </li> </ul> </li> </ul> |

## tokenmt(7ipp)

- o Else, outcome is RED.

Note that if `peak_burst` and `yellow_next_actions` are not specified (that is, a single rate meter with two outcomes), the outcome is never YELLOW.

When operating as a Two Rate meter, the outcome (level) is decided as follows:

- Update tokens in C and P
  - o Compute no. of tokens accumulated since the last time a packet was seen at the committed and peak rates as
$$T_c(t) = \text{committed rate} * (t - t')$$
$$T_p(t) = \text{peak rate} * (t - t')$$
(where  $t'$  is the time the last packet was seen)
  - o Add  $T_c$  to C up to a maximum of committed burst size
  - o Add  $T_p$  to P up to a maximum of peak burst size
- Decide outcome
  - o If not color aware
    - o If  $B > P$ , outcome is RED.
    - o Else, if  $B > C$ , outcome is YELLOW and  $P == B$
    - o Else, outcome is GREEN and  $C == B \ \& \ P == B$
  - o Else,
    - o obtain DSCP from packet
    - o obtain color from `color_map`, `color_map[DSCP]`
    - o if (color is RED) or ( $B > P$ ), outcome is RED
    - o Else, if (color is YELLOW) or ( $B > C$ ), outcome is YELLOW and  $P == B$
    - o Else, outcome is GREEN and  $C == B \ \& \ P == B$

**STATISTICS** The tokenmt module exports the following statistics through `kstat`:

Global statistics:

```
module: tokenmt instance: <action id>
 name: tokenmt statistics class <action name>
 epackets <number of packets in error>
 green_bits <number of bits in green>
 green_packets <number of packets in green>
 red_bits <number of bits in red>
 red_packets <number of packets in red>
 yellow_bits <number of bits in yellow>
 yellow_packets <number of packets in yellow>
```

**FILES** /kernel/ipp/sparcv9/tokenmt  
64-bit module (SPARC only.)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability   | SUNWqos         |

**SEE ALSO** `ipqosconf(1M)`, `dlcosmk(7IPP)`, `dscpmk(7IPP)`, `flowacct(7IPP)`, `ipqos(7IPP)`, `ipgpc(7IPP)`, `tswtclmt(7IPP)`

*RFC 2697, A Single Rate Three Color Marker* J. Heinanen, R. Guerin — The Internet Society, 1999

tokenmt(7ipp)

*RFC 2698, A Two Rate Three Color Marker* J. Heinanen, R. Guerin — The Internet Society, 1999

tpf(7D)

**NAME** tpf – Platform Specific Module (PSM) for Tricord Systems Enterprise Server Models ES3000, ES4000 and ES5000.

**DESCRIPTION** tpf provides the platform dependent functions for Solaris x86 MP support. These functions adhere to the PSMI Specifications. (Platform Specific Module Interface Specifications.) Tricord Systems Enterprise Servers are Intel APIC based MP platforms which run from 1 to 12 Intel processors. The tpf psm supports dynamic interrupt distribution across all processors in an MP configuration.

The psm is automatically invoked on an ESxxxx platform at system boot time.

**FILES** /kernel/mach/tpf MP module.

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Architecture   | x86             |

**SEE ALSO** attributes(5)

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>               | tsalarm – Alarm device driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>SYNOPSIS</b>           | tsalarm@0:ctl                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>DESCRIPTION</b>        | The <code>tsalarm</code> driver is a Multi-threaded, loadable non-STREAMS pseudo driver that manages ALOM alarms. The <code>tsalarm</code> driver provides an interface through which alarm relays can be controlled on SUNW,Netra-240 and SUNW,Netra-440 platforms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>HARDWARE INTERFACE</b> | The alarm hardware differs depending on platform. The Netra 240 and 440 platforms features four dry contact alarm relays which are controlled by ALOM. You can set each alarm to “on” or “off” by using <code>ioctl</code> interfaces provided from the host. The four alarms are labeled as “critical,” “major,” “minor,” and “user.” The user alarm is set by a user application depending on system condition. LED’s in front of the box provide a visual indication of the four alarms. The number of alarms and their meanings/labels may vary across platforms.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>IOCTLS</b>             | <p>The interface provided by the <code>tsalarm</code> driver comprises <code>ioctls</code> that enable applications to manipulate the alarm module. The alarm module is accessed via two device nodes: <code>i) /dev/lom</code> and <code>/dev/tsalarm:ctl</code>.</p> <p>The following <code>ioctls</code> are supported by the <code>/dev/lom</code> and <code>/dev/tsalarm:ctl</code> devices:</p> <p><b>TSIOCALCTL</b> - Turn an alarm on or off.<br/> The argument is a pointer to the <code>ts_aldata_t/lom_aldata_t</code> structure. This structure is described below. <code>alarm_no</code> member is an integer which specifies the alarm to which the command is to be applied. The <code>alarm_state/state</code> structure member indicates the state to which the alarm should be set (where 0 == off). An error (EINVAL) is returned if either an invalid <code>alarm_no</code> or invalid <code>alarm_state</code> is provided.</p> <p><b>TSIOCALSTATE</b> - Get the state of the alarms.<br/> The argument is a pointer to the <code>ts_aldata_t/lom_aldata_t</code> structure. This structure is described below. <code>alarm_no</code> member is an integer which indicates the alarm to which the command will be applied. The <code>alarm_state</code> member holds the alarm’s current state and is filled in by the driver. A zero indicates that the alarm is off. An error (EINVAL) is returned if an invalid <code>alarm_no</code> is provided. The structures and definitions for the values are defined below.</p> <p>Alarm values:</p> <p>The following old style values are defined in <code>&lt;lom.io.h&gt;</code></p> <pre> #define ALARM_NUM_0      0 /* number of zero'th alarm */ #define ALARM_NUM_1      1 /* number of first alarm */ #define ALARM_NUM_2      2 /* number of second alarm */ #define ALARM_NUM_3      3 /* number of third alarm */ </pre> <p>Alarm values defined in <code>&lt;lom.io.h&gt;</code></p> |

## tsalarm(7D)

```
#define ALARM_OFF 0 /* Turn off alarm */
#define ALARM_ON 1 /* Turn on alarm */
```

### Alarm Data Structure:

This structure is defined in `<lom.io.h>`

```
typedef struct {
 int alarm_no; /* alarm to apply command to */
 int alarm_state; /* state of alarm (0 == off) */
} ts_aldata_t;
```

Use the following LOM interfaces to get and set the alarms. These definitions are included in `<lom_io.h>`

```
#define ALARM_CRITICAL 0 /* number of critical alarm */
#define ALARM_MAJOR 1 /* number of major alarm */
#define ALARM_MINOR 2 /* number of minor alarm */
#define ALARM_USER 3 /* number of user alarm */
```

The following alarm data structure is provided in `<lom_io.h>`:

```
typedef struct {
 int alarm_no;
 int state;
} lom_aldata_t;
```

### ERRORS

An `open()` will fail if:

ENXIO

The driver is not installed in the system.

An `ioctl()` will fail if:

EFAULT

There was a hardware failure during the specified operation.

EINVAL

The alarm number specified is not valid or an invalid value was supplied.

ENXIO

The driver is not installed in the system or the monitor callback routine could not be scheduled.

### EXAMPLES

How to set an alarm:

```
#include <sys/unistd.h>
#include <fcntl.h>
#include <stdio.h>
```

```

#include <lom_io.h>

#define LOM_DEVICE "/dev/lom"

int
main()
{
 lom_aldata_t lld;
 int fd = open(LOM_DEVICE, O_RDWR);

 if (fd == -1) {
 printf("Error opening device: %s\n", LOM_DEVICE);
 exit (1);
 }

 lld.alarm_no = ALARM_CRITICAL; /* Set the critical alarm */
 lld.state = ALARM_ON; /* Set the alarm */

 if (ioctl(fd, LOMIOCALCTL, (char *)&lld) != 0)
 printf("Setting alarm failed");
 else
 printf("Alarm set successfully");

 close(fd);
}

```

**FILES**

/dev/lom  
LOM device.

/dev/tsalarm:ctl  
Alarm control device.

/platform/platform/kernel/drv/sparcv9/tsalarm  
Device driver module.

/platform/SUNW,Netra-240/kernel/drv/tsalarm.conf  
Driver configuration file.

**ATTRIBUTES**

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability   | SUNWcarx.u      |
| Architecture   | SPARC           |

**SEE ALSO**

[attributes\(5\)](#)  
*Writing Device Drivers*

tswtclmt(7ipp)

| <b>NAME</b>        | tswtclmt – Time Sliding Window Three Conformance Level Meter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                |                 |              |         |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-----------------|--------------|---------|
| <b>DESCRIPTION</b> | <p>The Time Sliding Window Three Conformance level meter (<code>tswtcl</code>) meters a traffic stream and determines the conformance level of its packets.</p> <p>Packets are deemed to belong to one of the three levels, Red, Yellow or Green, depending on the committed and peak rate.</p> <p>The meter provides an estimate of the running average bandwidth. It takes into account burstiness and smoothes out its estimate to approximate the longer-term measured sending rate of the traffic stream.</p> <p>The estimated bandwidth approximates the running average bandwidth of the traffic stream over a specific window (time interval). <code>tswtcl</code> estimates the average bandwidth using a time-based estimator. When a packet arrives for a class, <code>tswtcl</code> re-computes the average rate by using the rate in the last window (time interval) and the size of the arriving packet. The window is then slid to start at the current time (the packet arrival time). If the computed rate is less than the committed configuration parameter, the packet is deemed Green; else if the rate is less than the peak rate, it is Yellow; else Red. To avoid dropping multiple packets within a TCP window, <code>tswtcl</code> probabilistically assigns one of the three conformance level to the packet.</p> |                |                 |              |         |
| <b>STATISTICS</b>  | <p>The <code>tswtcl</code> module exports global and per-class statistics through <code>kstat</code>:</p> <p>Global statistics:</p> <pre>module: tswtclmt                               instance: &lt;action id&gt;   name: tswtclmt statistics                       class &lt;action name&gt;     green_bits                                   &lt;number of bit in green&gt;     green_packets                                &lt;number of packets in green&gt;     red_bits                                     &lt;number of bits in red&gt;     red_packets                                  &lt;number of packets in red&gt;     yellow_bits                                  &lt;number of bits in yellow&gt;     yellow_packets                              &lt;number of packets in yellow&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                |                 |              |         |
| <b>FILES</b>       | <code>/kernel/ipp/sparcv9/tswtclmt</code><br>64-bit module (SPARC only.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                |                 |              |         |
| <b>ATTRIBUTES</b>  | See <code>attributes(5)</code> for descriptions of the following attributes:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                |                 |              |         |
|                    | <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWqos</td></tr></tbody></table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Availability | SUNWqos |
| ATTRIBUTE TYPE     | ATTRIBUTE VALUE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                |                 |              |         |
| Availability       | SUNWqos                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                |                 |              |         |
| <b>SEE ALSO</b>    | <p><code>ipqosconf(1M)</code>, <code>dlcosmk(7IPP)</code>, <code>dscpmk(7IPP)</code>, <code>flowacct(7IPP)</code>, <code>ipqos(7IPP)</code>, <code>ipgpc(7IPP)</code>, <code>tokenmt(7IPP)</code></p> <p><i>RFC 2859, A Time Sliding Window Three Colour Marker (TSWTCM)</i> W. Fang, N. Seddigh, B. Nandy — The Internet Society, 2000</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                |                 |              |         |



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | ttcompat – V7, 4BSD and XENIX STREAMS compatibility module                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>SYNOPSIS</b>    | <pre>#define BSD_COMP #include &lt;sys/stropts.h&gt; #include &lt;sys/ioctl.h&gt; ioctl(<i>fd</i>, I_PUSH, "ttcompat");</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>DESCRIPTION</b> | <p>ttcompat is a STREAMS module that translates the <code>ioctl</code> calls supported by the older Version 7, 4BSD, and XENIX terminal drivers into the <code>ioctl</code> calls supported by the <code>termio</code> interface (see <a href="#">termio(7I)</a>). All other messages pass through this module unchanged; the behavior of <code>read</code> and <code>write</code> calls is unchanged, as is the behavior of <code>ioctl</code> calls other than the ones supported by <code>ttcompat</code>.</p> <p>This module can be automatically pushed onto a stream using the <code>autopush</code> mechanism when a terminal device is opened; it does not have to be explicitly pushed onto a stream. This module requires that the <code>termios</code> interface be supported by the modules and the application can push the driver downstream. The <code>TCGETS</code>, <code>TCSETS</code>, and <code>TCSETSF</code> <code>ioctl</code> calls must be supported. If any information set or fetched by those <code>ioctl</code> calls is not supported by the modules and driver downstream, some of the V7/4BSD/XENIX functions may not be supported. For example, if the CBAUD bits in the <code>c_cflag</code> field are not supported, the functions provided by the <code>sg_ispeed</code> and <code>sg_ospeed</code> fields of the <code>sgttyb</code> structure (see below) will not be supported. If the <code>TCFLSH</code> <code>ioctl</code> is not supported, the function provided by the <code>TIOCFLUSH</code> <code>ioctl</code> will not be supported. If the <code>TCXONC</code> <code>ioctl</code> is not supported, the functions provided by the <code>TIOCSTOP</code> and <code>TIOCSTART</code> <code>ioctl</code> calls will not be supported. If the <code>TIOCMBIS</code> and <code>TIOCMBIC</code> <code>ioctl</code> calls are not supported, the functions provided by the <code>TIOCSDTR</code> and <code>TIOCCDTR</code> <code>ioctl</code> calls will not be supported.</p> <p>The basic <code>ioctl</code> calls use the <code>sgttyb</code> structure defined by <code>&lt;sys/ttold.h&gt;</code> (included by <code>&lt;sys/ioctl.h&gt;</code>):</p> <pre>struct sgttyb {     char    sg_ispeed;     char    sg_ospeed;     char    sg_erase;     char    sg_kill;     int     sg_flags; };</pre> <p>The <code>sg_ispeed</code> and <code>sg_ospeed</code> fields describe the input and output speeds of the device. If the speed set on the device is over B38400, then it is reported as B38400 for compatibility reasons. If it is set to B38400 and the current speed is over B38400, the change is ignored. See <code>TIOCGETP</code> and <code>TIOCSETP</code> below. The <code>sg_erase</code> and <code>sg_kill</code> fields of the argument structure specify the erase and kill characters respectively, and reflect the values in the <code>VERASE</code> and <code>VKILL</code> members of the <code>c_cc</code> field of the <code>termios</code> structure.</p> <p>The <code>sg_flags</code> field of the argument structure contains several flags that determine the system's treatment of the terminal. They are mapped into flags in fields of the terminal state, represented by the <code>termios</code> structure.</p> |

ttcompat(7M)

Delay type 0 (NL0, TAB0, CR0, FF0, BS0) is always mapped into the equivalent delay type 0 in the `c_oflag` field of the `termios` structure. Other delay mappings are performed as follows:

| <code>sg_flags</code> | <code>c_oflag</code> |
|-----------------------|----------------------|
| BS1                   | BS1                  |
| FF1                   | VT1                  |
| CR1                   | CR2                  |
| CR2                   | CR3                  |
| CR3                   | CR0 (not supported)  |
| TAB1                  | TAB1                 |
| TAB2                  | TAB2                 |
| XTABS                 | TAB3                 |
| NL1                   | ONLRET CR1           |
| NL2                   | NL1                  |
| NL3                   | NL0 (not supported)  |

If previous `TIOCLSET` or `TIOCLBIS` `ioctl` calls have not selected `LITOUT` or `PASS8` mode, and if `RAW` mode is not selected, the `ISTRIP` flag is set in the `c_iflag` field of the `termios` structure, and the `EVENP` and `ODDP` flags control the parity of characters sent to the terminal and accepted from the terminal, as follows:

|                                                       |                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 0 (neither <code>EVENP</code> nor <code>ODDP</code> ) | Parity is not to be generated on output or checked on input. The character size is set to <code>CS8</code> and the <code>PARENB</code> flag is cleared in the <code>c_cflag</code> field of the <code>termios</code> structure.                                                                                                                        |
| <code>EVENP</code>                                    | Even parity characters are to be generated on output and accepted on input. The <code>INPCK</code> flag is set in the <code>c_iflag</code> field of the <code>termios</code> structure, the character size is set to <code>CS7</code> and the <code>PARENB</code> flag is set in the <code>c_iflag</code> field of the <code>termios</code> structure. |
| <code>ODDP</code>                                     | Odd parity characters are to be generated on output and accepted on input. The <code>INPCK</code> flag is set in the <code>c_iflag</code> , the character size is set to <code>CS7</code> and the <code>PARENB</code> and <code>PARODD</code> flags are set in the <code>c_iflag</code> field of the <code>termios</code> structure.                   |

EVENP | ODDP or ANYP

Even parity characters are to be generated on output and characters of either parity are to be accepted on input. The INPCK flag is cleared in the `c_iflag` field, the character size is set to CS7 and the PARENB flag is set in the `c_iflag` field of the `termios` structure.

The RAW flag disables all output processing (the OPOST flag in the `c_oflag` field, and the XCASE and IEXTEN flags in the `c_iflag` field are cleared in the `termios` structure) and input processing (all flags in the `c_iflag` field other than the IXOFF and IXANY flags are cleared in the `termios` structure). Eight bits of data, with no parity bit are accepted on input and generated on output; the character size is set to CS8 and the PARENB and PARODD flags are cleared in the `c_cflag` field of the `termios` structure. The signal-generating and line-editing control characters are disabled by clearing the ISIG and ICANON flags in the `c_iflag` field of the `termios` structure.

The CRMOD flag turns input carriage return characters into linefeed characters, and output linefeed characters to be sent as a carriage return followed by a linefeed. The ICRNL flag in the `c_iflag` field, and the OPOST and ONLCR flags in the `c_oflag` field, are set in the `termios` structure.

The LCASE flag maps upper-case letters in the ASCII character set to their lower-case equivalents on input (the IUCLC flag is set in the `c_iflag` field), and maps lower-case letters in the ASCII character set to their upper-case equivalents on output (the OLCUC flag is set in the `c_oflag` field). Escape sequences are accepted on input, and generated on output, to handle certain ASCII characters not supported by older terminals (the XCASE flag is set in the `c_lflag` field).

Other flags are directly mapped to flags in the `termios` structure:

| sg_flags | Flags in termios structure                         |
|----------|----------------------------------------------------|
| CBREAK   | Complement of ICANON in <code>c_lflag</code> field |
| ECHO     | ECHO in <code>c_lflag</code> field                 |
| TANDEM   | IXOFF in <code>c_iflag</code> field                |

Another structure associated with each terminal specifies characters that are special in both the old Version 7 and the newer 4BSD terminal interfaces. The following structure is defined by `<sys/ttold.h>`:

```
struct tchars {
 char t_intrc; /* interrupt */
 char t_quitc; /* quit */
 char t_startc; /* start output */
 char t_stopc; /* stop output */
 char t_eofc; /* end-of-file */
 char t_brkc; /* input delimiter (like nl) */
}
```

```
};
```

XENIX defines the `tchar` structure as `tc`. The characters are mapped to members of the `c_cc` field of the `termios` structure as follows:

```
tchars c_cc index
t_intrc VINTR
t_quitc VQUIT
t_startc VSTART
t_stopc VSTOP
t_eofc VEOF
t_brkc VEOL
```

Also associated with each terminal is a local flag word (`TIOCLSET` and `TIOCLGET`), specifying flags supported by the new 4BSD terminal interface. Most of these flags are directly mapped to flags in the `termios` structure:

| Local flags | Flags in termios structure                            |
|-------------|-------------------------------------------------------|
| LCRTBS      | Not supported                                         |
| LPRTERA     | ECHOPRT in the <code>c_lflag</code> field             |
| LCRTERA     | ECHOE in the <code>c_lflag</code> field               |
| LTILDE      | Not supported                                         |
| LMDMBUF     | Not supported                                         |
| LTOSTOP     | TOSTOP in the <code>c_lflag</code> field              |
| LFLUSHO     | FLUSHO in the <code>c_lflag</code> field              |
| LNOHANG     | CLOCAL in the <code>c_cflag</code> field              |
| LCRTKIL     | ECHOKE in the <code>c_lflag</code> field              |
| LPASS8      | CS8 in the <code>c_cflag</code> field                 |
| LCTLECH     | CTLECH in the <code>c_lflag</code> field              |
| LPENDIN     | PENDIN in the <code>c_lflag</code> field              |
| LDECCTQ     | Complement of IXANY in the <code>c_iflag</code> field |
| LNOFLSH     | NOFLSH in the <code>c_lflag</code> field              |

Each flag has a corresponding equivalent `sg_flags` value. The `sg_flags` definitions omit the leading "L"; for example, `TIOCSETP` with `sg_flags` set to `TOSTOP` is equivalent to `TIOCLSET` with `LTOSTOP`.

Another structure associated with each terminal is the `ltchars` structure which defines control characters for the new 4BSD terminal interface. Its structure is:

```

struct ltchars {
 char t_suspc; /* stop process signal */
 char t_dsuspc; /* delayed stop process signal */
 char t_rprntc; /* reprint line */
 char t_flushc; /*flush output (toggles) */
 char t_werasc; /* word erase */
 char t_lnextc; /* literal next character */
};

```

The characters are mapped to members of the `c_cc` field of the `termios` structure as follows:

| ltchars  | c_cc index |
|----------|------------|
| t_suspc  | VSUS       |
| t_dsuspc | VDSUSP     |
| t_rprntc | VREPRINT   |
| t_flushc | VDISCARD   |
| t_werasc | VWERASE    |
| t_lnextc | VLNEXT     |

**IOCTLS** ttcompat responds to the following `ioctl` calls. All others are passed to the module below.

- TIOCGETP                      The argument is a pointer to an `sgttyb` structure. The current terminal state is fetched; the appropriate characters in the terminal state are stored in that structure, as are the input and output speeds. If the speed is over B38400, then B38400 is returned. The values of the flags in the `sg_flags` field are derived from the flags in the terminal state and stored in the structure.
- TIOCEXCL                      Set "exclusive-use" mode; no further opens are permitted until the file has been closed.
- TIOCNXCL                      Turn off "exclusive-use" mode.
- TIOCSETP                      The argument is a pointer to an `sgttyb` structure. The appropriate characters and input and output speeds in the terminal state are set from the values in that structure, and the flags in the terminal state are set to match the values of the flags in the `sg_flags` field of that structure. The state is changed with a `TCSETS` `ioctl` so

## ttcompat(7M)

|           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | that the interface delays until output is quiescent, then throws away any unread characters, before changing the modes. If the current device speed is over B38400 for either input or output speed, and B38400 is specified through this interface for that speed, the actual device speed is not changed. If the device speed is B38400 or lower or if some speed other than B38400 is specified, then the actual speed specified is set.                                                                             |
| TIOCSETN  | The argument is a pointer to an <code>sgttyb</code> structure. The terminal state is changed as <code>TIOCSETP</code> would change it, but a <code>TCSETS</code> <code>ioctl</code> is used, so that the interface neither delays nor discards input.                                                                                                                                                                                                                                                                   |
| TIOCHPCL  | The argument is ignored. The <code>HUPCL</code> flag is set in the <code>c_cflag</code> word of the terminal state.                                                                                                                                                                                                                                                                                                                                                                                                     |
| TIOCFLUSH | The argument is a pointer to an <code>int</code> variable. If its value is zero, all characters waiting in input or output queues are flushed. Otherwise, the value of the <code>int</code> is treated as the logical OR of the <code>FREAD</code> and <code>FWRITE</code> flags defined by <code>&lt;sys/file.h&gt;</code> . If the <code>FREAD</code> bit is set, all characters waiting in input queues are flushed, and if the <code>FWRITE</code> bit is set, all characters waiting in output queues are flushed. |
| TIOCSBRK  | The argument is ignored. The break bit is set for the device. (This is not supported by <code>ttcompat</code> . The underlying driver must support <code>TIOCSBRK</code> .)                                                                                                                                                                                                                                                                                                                                             |
| TIOCCBRK  | The argument is ignored. The break bit is cleared for the device. (This is not supported by <code>ttcompat</code> . The underlying driver must support <code>TIOCCBRK</code> .)                                                                                                                                                                                                                                                                                                                                         |
| TIOCSDTR  | The argument is ignored. The Data Terminal Ready bit is set for the device.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| TIOCCDTR  | The argument is ignored. The Data Terminal Ready bit is cleared for the device.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| TIOCSTOP  | The argument is ignored. Output is stopped as if the <code>STOP</code> character had been typed.                                                                                                                                                                                                                                                                                                                                                                                                                        |

|            |                                                                                                                                                                                                                                                                                                                                                                                                |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TIOCPSTART | The argument is ignored. Output is restarted as if the START character had been typed.                                                                                                                                                                                                                                                                                                         |
| TIOCGETC   | The argument is a pointer to a <code>tchars</code> structure. The current terminal state is fetched, and the appropriate characters in the terminal state are stored in that structure.                                                                                                                                                                                                        |
| TIOCSETC   | The argument is a pointer to a <code>tchars</code> structure. The values of the appropriate characters in the terminal state are set from the characters in that structure.                                                                                                                                                                                                                    |
| TIOCLGET   | The argument is a pointer to an <code>int</code> . The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state and stored in the <code>int</code> pointed to by the argument.                                                                                                                                                    |
| TIOCLBIS   | The argument is a pointer to an <code>int</code> whose value is a mask containing flags to be set in the local flags word. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state; the specified flags are set, and the flags in the terminal state are set to match the new value of the local flags word.                 |
| TIOCLBIC   | The argument is a pointer to an <code>int</code> whose value is a mask containing flags to be cleared in the local flags word. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state; the specified flags are cleared, and the flags in the terminal state are set to match the new value of the local flags word.         |
| TIOCLSET   | The argument is a pointer to an <code>int</code> containing a new set of local flags. The flags in the terminal state are set to match the new value of the local flags word. (This <code>ioctl</code> was added because <code>sg_flags</code> was once a 16 bit value. The local modes controlled by TIOCLSET are equivalent to the modes controlled by TIOCSETP and <code>sg_flags</code> .) |

## ttcompat(7M)

|          |                                                                                                                                                                                                |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TIOCGLTC | The argument is a pointer to an <code>ltchars</code> structure. The values of the appropriate characters in the terminal state are stored in that structure.                                   |
| TIOCSLTC | The argument is a pointer to an <code>ltchars</code> structure. The values of the appropriate characters in the terminal state are set from the characters in that structure.                  |
| FIORDCHK | Returns the number of immediately readable characters. The argument is ignored. (This <code>ioctl</code> is handled in the stream head, not in the <code>ttcompat</code> module.)              |
| FIONREAD | Returns the number of immediately readable characters in the int pointed to by the argument. (This <code>ioctl</code> is handled in the stream head, not in the <code>ttcompat</code> module.) |

The following `ioctls` are returned as successful for the sake of compatibility. However, nothing significant is done (that is, the state of the terminal is not changed in any way, and no message is passed through to the underlying `tty` driver).

```
TIOCSETD LDOPEN
TIOCGETD LDCLOSE
DIOCSETP LDCHG
DIOCSETP LDSETP
DIOCGETP LDGETT
```

The following old `ioctls` are not supported by `ttcompat`, but are supported by Solaris `tty` drivers. As with all `ioctl` not otherwise listed in this documentation, these are passed through to the underlying driver and are handled there.

```
TIOCREMOTE
TIOCGWINSZ TIOCSWINSZ
```

The following `ioctls` are not supported by `ttcompat`, and are generally not supported by Solaris `tty` drivers. They are passed through, and the `tty` drivers return `EINVAL`.

```
LDSMAP TIOCNOTTY
LDGMAP TIOCOUTQ
LDNMAP
```

(Note: `LDSMAP`, `LDGMAP`, and `LDNMAP` are defined in `<sys/termios.h>`.)

**SEE ALSO** `ioctl(2)`, `termios(3C)`, `ldterm(7M)`, `termio(7I)`



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | tty – controlling terminal interface                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>DESCRIPTION</b> | The file <code>/dev/tty</code> is, in each process, a synonym for the control terminal associated with the process group of that process, if any. It is useful for programs or shell sequences that wish to be sure of writing messages on the terminal no matter how output has been redirected. It can also be used for programs that demand the name of a file for output, when typed output is desired and it is tiresome to find out what terminal is currently in use. |
| <b>FILES</b>       | <code>/dev/tty</code><br><code>/dev/tty*</code>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>SEE ALSO</b>    | <code>ports(1M)</code> , <code>console(7D)</code>                                                                                                                                                                                                                                                                                                                                                                                                                            |

## ttymux(7D)

| <b>NAME</b>        | ttymux – Serial I/O multiplexing STREAMS device driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                |                 |              |                             |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-----------------|--------------|-----------------------------|
| <b>SYNOPSIS</b>    | multiplexer@0,0:input<br>multiplexer@0,0:output                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                |                 |              |                             |
| <b>DESCRIPTION</b> | <p>ttymux is a STREAMS multiplexer driver that connects multiple serial devices to the system console. Using this driver, input from multiple physical devices can be multiplexed onto a single input stream for the system console. Output written to the console can be distributed to multiple physical devices to provide redundant console interfaces to a system. Input and output can be multiplexed to or from a separate list of devices.</p> <p>ttymux is a STREAM's multiplexer for serial drivers (such as <a href="#">se(7D)</a>) that comply with the Solaris terminal subsystem interface.</p> <p>Currently, multiplexer interfaces are provided for system console I/O only and not for general serial I/O multiplexing. Multiplexer interfaces are currently not available for all platforms. Please see NOTES.</p>                                                                                                                                                                                                                       |                |                 |              |                             |
| <b>FILES</b>       | /kernel/drv/sparcv9/ttymux 64-bit ELF kernel module                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                |                 |              |                             |
| <b>ATTRIBUTES</b>  | See <a href="#">attributes(5)</a> for descriptions of the following attributes:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                |                 |              |                             |
|                    | <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>SPARC (NetraCT series only)</td></tr></tbody></table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Architecture | SPARC (NetraCT series only) |
| ATTRIBUTE TYPE     | ATTRIBUTE VALUE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                |                 |              |                             |
| Architecture       | SPARC (NetraCT series only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                |                 |              |                             |
| <b>SEE ALSO</b>    | <a href="#">se(7D)</a> , <a href="#">termio(7I)</a><br><br><i>Writing Device Drivers</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                |                 |              |                             |
| <b>NOTES</b>       | <p>Successful loading of this driver and its services depends on the EEPROM or NVRAM settings in effect at the most recent system reboot. Without the platform firmware support, this feature cannot be enabled. Currently, this support is provided only on a NetraCT product family.</p> <p>Use caution when enabling this feature to perform console input multiplexing, particularly during super-user login. Because no security measures are enabled when the driver is in operation, you must clearly understand the security implications involved in using this feature and take appropriate measures to provide maximum protection to the host. This can include such steps as enabling input to physically secured console devices only.</p> <p>The ttymux driver does not handle the behavioral differences in control characteristics of different terminal types (for example, an ESCAPE sequence.) As a result, multiple terminal types are not supported simultaneously. Please refer to the platform user guide for more information.</p> |                |                 |              |                             |

|                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                               | tun, TUN, 6to4, 6to4tun, atun – tunneling STREAMS module                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>SYNOPSIS</b>                           | <pre>strmod/tun strmod/atun strmod/6to4tun</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>DESCRIPTION</b>                        | <p>tun, atun and 6to4tun are STREAMS modules that implement an IP-in-IP tunneling mechanism. IPv6-in-IPv4, IPv4-in-IPv4, IPv4-in-IPv6 and IPv6-in-IPv6 tunnels are supported.</p> <p>Tunnels are configured as point-to-point interfaces. Ipv4-in-Ipv4 allows IPv4 packets to be encapsulated within IPv4 packets. IPv6-in-IPv4 tunnels allow IPv6 packets to be encapsulated within IPv4 packets. IPv4-in-IPv6 tunnels allow IPv4 packets to be encapsulated within IPv6 packets. IPv6-in-IPv6 tunnels allow IPv6 packets to be encapsulated within IPv6 packets. Both the tunnel source and the tunnel destination are required to configure these type of tunnels. Configured tunnels support encapsulated multicast packets. See <code>ifconfig(1M)</code> for examples of these tunnel configurations.</p> <p>The <code>atun</code> module is used to configure automatic tunnels. It supports IPv6 packets encapsulated within IPv4 packets. An IPv4 address is required for the tunnel source of these interfaces and the IPv4 compatible IPv6 source address must match this address. IPv6 packets using this interface must have IPv4 compatible source and destination addresses. Automatic tunnels are not point-to-point, and they do not allow multicast packets to be sent. If the destination of an automatic tunnel is a router, the packets will not be forwarded.</p> <p>The <code>6to4tun</code> module is used to configure 6to4 tunnels as described in <i>RFC 3056</i>. It implements automatic tunneling of IPv6 within IPv4. This IPv6 transition mechanism allows isolated IPv6 networks or nodes attached to an IPv4 network to communicate with other IPv6 networks with little configuration. See <code>ifconfig(1M)</code> for an example of how to configure a 6to4 tunnel.</p> <ul style="list-style-type: none"> <li>■ Network startup scripts look at <code>/etc/hostname.ip.*</code>, <code>/etc/hostname6.ip.*</code>, <code>/etc/hostname.ip6.*</code> and <code>/etc/hostname6.ip6.*</code> to find the available tunneling interfaces.</li> <li>■ The same tunnel source address (<code>tsrc</code>) and destination address (<code>tdst</code>) is be used for all instances (<code>luns</code>) of a specific interface.</li> <li>■ Tunnels do not support snooping. Instead, a filter made up of the combination of addresses can be used on the physical interface to capture relevant packets.</li> <li>■ If there is a tunnel set up between two multicast routers, then multicast routing should be configured to use the tunnel, rather than a special multicast routing virtual interface.</li> </ul> <p>The tunnel module is architected to be plumbed between two instances of IP.</p> <p>The following <code>ioctl()</code> calls may be used to configure a tunneling interface. The <code>ioctl()</code>s are defined in <code>&lt;sys/socket.h&gt;</code>. This structure is defined in <code>&lt;net/if.h&gt;</code>.</p> |
| <b>APPLICATION PROGRAMMING INTERFACES</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## tun(7M)

```
/* currently tunnels only support IPv4 or IPv6 */
enum ifta_proto {
 IFTAP_INVALID,
 IFTAP_IPV4,
 IFTAP_IPV6
};

#define IFTUN_SECINFOLEN 8
#define IFTUN_VERSION 1

/* tunnel configuration structure */
struct iftun_req {
 char ifta_lifr_name[LIFNAMSIZ]; /* if name */
 struct sockaddr_storage ifta_saddr; /* source address */
 struct sockaddr_storage ifta_daddr; /* destination address */
 uint_t ifta_flags; /* See below */
 /* IP version information is read only */
 enum ifta_proto ifta_upper; /* IP version above tunnel */
 enum ifta_proto ifta_lower; /* IP version below tunnel */
 uint_t ifta_vers; /* Version number */
 uint32_t ifta_secinfo[IFTUN_SECINFOLEN]; /* Security prefs. */
 int16_t ifta_encap_lim; /* Encapsulation limit */
 uint8_t ifta_hop_limit; /* Hop limit */
};

/* These flags are set to indicate which members are valid */

#define IFTUN_SRC 0x01
#define IFTUN_DST 0x02
#define IFTUN_SECURITY 0x04
#define IFTUN_ENCAP 0x08
#define IFTUN_HOPLIMIT 0x10
```

The `ifta_vers` field indicates what IPsec request structure is overlaid on top of `ifta_secinfo`. The current value of `IFTUN_VERSION` implies an overlay of `ipsec_req_t`. See [ipsec\(7P\)](#).

### SIOCSTUNPARAM

Set tunnel parameters. This `ioctl()` allows the tunnel's source or destination address or hop limit or encapsulation limit to be set. The `IFTUN_SRC` bit set in `ta_flags` indicates that the tunnel should bound to the source address supplied in `ta_saddr`. The source must be a valid configured interface IP address. The `IFTUN_DST` bit set in `ta_flags` indicates that the tunnel should bound to the destination address supplied in `ta_daddr`. The destination address must be reachable. The `IFTUN_ENCAP` bit set in `ifta_flags` indicates that the tunnel's encapsulation limit should be set to the value supplied in `ifta_encap_lim`. The encapsulation limit is valid only for IPv4-in-IPv6 and IPv6-in-IPv6 tunnels. Valid encapsulation limit values are 0 through 255. Negative values indicate that no encapsulation limit is desired. The `IFTUN_HOPLIMIT` bit set in `ta_flags` indicates that the tunnel's hop limit should be set to the value supplied in `ifta_hop_limit`. In the case of IPv4-in-IPv4 and IPv6-in-IPv4 tunnels, the hop limit is placed in the IPv4 header's TTL field. In the case of IPv4-in-IPv6 and IPv6-in-IPv6 tunnels, the hop limit is placed in the IPv6 header's hop limit field.

**SIOCGTUNPARAM**

Get tunnel parameters. Valid fields are indicated by the returned value of `ta_flags` bitmask. The version of IP plumbed above or below the tunnel may be determined by inspecting `ta_upper` and `ta_lower` by comparing the members against the mutually exclusive defined values `IFTAP_INVALID`, `IFTAP_IPV4`, and `IFTAP_IPV6`.

**Tunnels and DLPI**

The tunnel module is a DLPI style 2 service provider. All `M_PROTO` and `M_PCPROTO` type messages are interpreted as DLPI primitives. Valid DLPI primitives are defined in `<sys/dlpi.h>`. Refer to [dlpi\(7P\)](#) for more information. An explicit `DL_ATTACH_REQ` message by the user is required to associate the opened stream with a particular device (`ppa`). The `ppa` indicates the corresponding device instance (`unit`) number. The device is initialized on first attach and deinitialized (stopped) on last detach.

The values returned by the module in the `DL_INFO_ACK` primitive in response to the `DL_INFO_REQ` from the user are as follows:

- The maximum SDU is usually 4196 ("`ip_max_mtu` - size of IP header").
- The minimum SDU is 1.
- The `dlsap` address length is 0 for configured tunnels and non-zero for automatic tunnels.
- The MAC type is `DL_OTHER`.
- The `sap` length value is 0.
- The service mode is `DL_CLDLS`.
- No optional quality of service (QOS) support is included at present so the QOS fields are 0.
- The provider style is `DL_STYLE2`.
- The version is `DL_VERSION_2`.
- The broadcast address value is 0

Once in the `DL_ATTACHED` state, the user must send a `DL_BIND_REQ` to associate a particular SAP (Service Access Pointer) with the stream. The tunneling module interprets the `sap` field within the `DL_BIND_REQ` as an IP "type" therefore the valid value for the `sap` field is `IP_DL_SAP`.

Once in the `DL_BOUND` state, the user may transmit packets through the tunnel by sending `DL_UNITDATA_REQ` messages to the tunnel module. Configured tunnels will encapsulate the packet with the appropriate IP header using the source and destination specified by `tsrc` and `tdst` parameters of `ifconfig(1M)`. The tunnel module will decapsulate received packets and route them to the first open and bound stream having a `sap`, `tsrc` and `tdst` which matches the the configured information. Packets are routed to exactly one open stream and not duplicated.

The module does not support additional primitives. `DL_ERROR_ACK` with the `dl_error` set to `DL_UNSUPPORTED` will be returned in the case that an unsupported DLPI primitive is encountered.

tun(7M)

**SECURITY  
CONSIDERATIONS**

A tunnel creates what appears to be a physical interface to IP. It can be "trusted" as a physical link only so far as the underlying security protocols, if used, can be trusted. If the security associations (see [ipsec\(7P\)](#)) are securely set up then the tunnel can be trusted in that packets that come off the tunnel came from the peer specified in the tunnel destination. If this trust exists, per-interface IP forwarding can be used to create a Virtual Private Network ("VPN"). See [ip\(7P\)](#).

**ATTRIBUTES**

See [attributes\(5\)](#) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Availability        | SUNWcsr         |
| Interface Stability | Evolving        |

**SEE ALSO**

[ifconfig\(1M\)](#), [attributes\(5\)](#), [ip\(7P\)](#), [ipsec\(7P\)](#)

*System Administration Guide: IP Services*

Gilligan, R. and Nordmark, E., *RFC 1933, Transition Mechanisms for IPv6 Hosts and Routers*, The Internet Society, 1996.

Conta, A. and Deering, S. *RFC 2473, Generic Packet Tunneling in IPv6 Specification*, The Internet Society, 1998.

Carpenter, B and Moore, K. *RFC 3056, Connection of IPv6 Domains via IPv4 Clouds*, The Internet Society, 2001

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |   |                                |   |                               |   |                                  |   |                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|--------------------------------|---|-------------------------------|---|----------------------------------|---|---------------------------------|
| <b>NAME</b>        | uata – IDE Host Bus Adapter Driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |   |                                |   |                               |   |                                  |   |                                 |
| <b>SYNOPSIS</b>    | <code>ide@unit-address</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |   |                                |   |                               |   |                                  |   |                                 |
| <b>DESCRIPTION</b> | <p>The uata host bus adapter driver is a nexus driver that supports the <code>ide</code> interface on SPARC platforms. The driver supports ultra DMA mode-4 (ATA66). The driver also attempts to set the disk and ATAPI CD-ROM drive to maximum supported speed for the device.</p> <p>Currently, the uata driver supports IDE controllers CMD646U and Acer Southbridge M5229.</p> <p>The uata driver supports two channels concurrently with two devices connected per channel. The devices are logically numbered from 0 to 3:</p> <table border="0"> <tr> <td style="padding-right: 20px;">0</td> <td>Master disk on primary channel</td> </tr> <tr> <td>1</td> <td>Slave disk on primary channel</td> </tr> <tr> <td>2</td> <td>Master disk on secondary channel</td> </tr> <tr> <td>3</td> <td>Slave disk on secondary channel</td> </tr> </table>                                                                                                                                                                                                                                                                           | 0 | Master disk on primary channel | 1 | Slave disk on primary channel | 2 | Master disk on secondary channel | 3 | Slave disk on secondary channel |
| 0                  | Master disk on primary channel                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |   |                                |   |                               |   |                                  |   |                                 |
| 1                  | Slave disk on primary channel                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |   |                                |   |                               |   |                                  |   |                                 |
| 2                  | Master disk on secondary channel                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |   |                                |   |                               |   |                                  |   |                                 |
| 3                  | Slave disk on secondary channel                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |   |                                |   |                               |   |                                  |   |                                 |
| <b>FILES</b>       | <code>/kernel/drv/uata</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |   |                                |   |                               |   |                                  |   |                                 |
| <b>SEE ALSO</b>    | <p><code>prtconf(1M)</code>, <code>driver.conf(4)</code>, <code>attributes(5)</code></p> <p><i>Writing Device Drivers</i></p> <p><i>X3T10 ATA-4 specifications</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |   |                                |   |                               |   |                                  |   |                                 |
| <b>DIAGNOSTICS</b> | <p>In addition to being logged, the following messages may appear on the system console:</p> <p><code>ddi_get_iblock_cookie failed</code><br/> The driver could not obtain the interrupt cookie; the attach may fail.</p> <p><code>Drive not ready before set_features</code><br/> Indicates a fatal problem; the drives are not ready to be programmed and features cannot be set. (During the driver initialization process, the driver must set the features for the drive, including <code>dma</code> and <code>pio</code>.)</p> <p><code>Error set after issuing Set Feature command</code><br/> Indicates a fatal problem; the driver's error bit was set after the set feature command was issued. (During the driver initialization process, the driver must set the features for the drive, including <code>dma</code> and <code>pio</code>.)</p> <p><code>Interrupt not seen after set_features</code><br/> Indicates a fatal problem with the drive; features could not be set.</p> <p><code>ata_controller - set_features failed</code><br/> Indicates a fatal problem with the drive; features could not be set.</p> |   |                                |   |                               |   |                                  |   |                                 |

## uata(7D)

? target <number> lun 0  
Displayed at boot up time to indicate that the target <number> was identified, where <number> is a decimal value.

resid  
Residual number of bytes in data transfer and the I/O operation could not be finished completely.

ghd\_timer\_newstate: HBA reset failed  
Generally indicates a fatal condition; I/O operation could not be completed following reset of the channel.

timeout: <message> chno =<number> target=<number>  
A timeout occurred because of <message> on device (target=<number>) on channel (chno =<number>). Where <message> could be either early abort, early timeout, abort request, abort device, reset target or reset bus.

ata\_controller - Drive not ready before command <number>  
The drive did not respond before issuing the command <number> to the controller; command <number> will not be issued to the drive. (<number> is the hexadecimal opcode for the sleep or standby commands, which are issued when the drive transitions between power management states.)

ata\_controller - Command <number> failed  
Command <number> failed on the drive. (<number> is the hexadecimal opcode for the sleep or standby commands, which are issued when the drive transitions between power management states.)

ata\_controller - Command <number> returned error  
Command <number> returned error. (<number> is the hexadecimal opcode for the sleep or standby commands, which are issued when the drive transitions between power management states.)

ata\_controller - Cannot take drive <number> to sleep  
The disk will not transition to sleep state. (Indicates that the driver could not set the device to sleep mode while performing power management functions.)

ata\_controller - Cannot reset secondary/primary channel  
The disk will not transition from sleep to active state.

ata\_controller - Unsupported Controller Vendor 0x13d0, Device 0x43f1, Revision 0x034  
An unsupported ata controller was found on the system and prints <ID>, device id and revision of the controller, where <ID> represents the hexadecimal vendor ID.

Changing the mode of targ: <number> to Ultra DMA mode: <number>  
For the timedout command, the driver attempts to recover by changing speed to lower values and retrying the command. This message indicates to which mode the driver is attempting to re-program the drive, where <number> is a decimal value.

Changing the mode of targ: <number> to Multi DMA mode: <number>  
For the timedout command, the driver attempts to recover by changing speed to lower values and retrying the command. This message indicates to which mode the



uata(7D)

driver is attempting to re-program the drive, where *<number>* is a decimal value.

These messages are informational and indicate that a timeout occurred for a I/O request. The uata driver recovers from these states automatically unless there is a fatal error.

## udfs(7FS)

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                  | udfs – universal disk format file system                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>DESCRIPTION</b>           | <p>The <code>udfs</code> file system is a file system type that allows user access to files on Universal Disk Format (UDF) disks from within the Solaris operating environment. Once mounted, a <code>udfs</code> file system provides standard Solaris file system operations and semantics. That is, users can read files, write files, and list files in a directory on a UDF device and applications can use standard UNIX system calls on these files and directories.</p> <p>Because <code>udfs</code> is a platform-independent file system, the same media can be written to and read from by any operating system or vendor.</p>                                                                                   |
| <b>Mounting File Systems</b> | <p><code>udfs</code> file systems are mounted using:</p> <pre>mount-F udfs -o rw/ro device-special</pre> <p>Use:</p> <pre>mount /udfs</pre> <p>if the <code>/udfs</code> and device special file <code>/dev/dsk/c0t6d0s0</code> are valid and the following line (or similar line) appears in your <code>/etc/vfstab</code> file:</p> <pre>/dev/dsk/c0t6d0s0 - /udfs udfs - no ro</pre> <p>The <code>udfs</code> file system provides read-only support for ROM, RAM, and sequentially-recordable media and read-write support on RAM media.</p> <p>The <code>udfs</code> file system also supports regular files, directories, and symbolic links, as well as device nodes such as block, character, FIFO, and Socket.</p> |
| <b>SEE ALSO</b>              | <code>mount(1M)</code> , <code>mount_udfs(1M)</code> , <code>vfstab(4)</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>NOTES</b>                 | <p>Invalid characters such as "NULL" and "/" and invalid file names such as "." and ".." will be translated according to the following rule:</p> <p>Replace the invalid character with an "_" then append the file name with # followed by a 4 digit hex representation of the 16-bit CRC of the original <code>FileIdentifier</code>. For example, the file name "." will become "__#4C05"</p>                                                                                                                                                                                                                                                                                                                             |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | udp, UDP – Internet User Datagram Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>SYNOPSIS</b>    | <pre>#include &lt;sys/socket.h&gt; #include &lt;netinet/in.h&gt;  s = socket(AF_INET, SOCK_DGRAM, 0); s = socket(AF_INET6, SOCK_DGRAM, 0); t = t_open("/dev/udp", O_RDWR); t = t_open("/dev/udp6", O_RDWR);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>DESCRIPTION</b> | <p>UDP is a simple datagram protocol which is layered directly above the Internet Protocol (“IP”) or the Internet Protocol Version 6 (“IPv6”). Programs may access UDP using the socket interface, where it supports the <code>SOCK_DGRAM</code> socket type, or using the Transport Level Interface (“TLI”), where it supports the connectionless (<code>T_CLTS</code>) service type.</p> <p>Within the socket interface, UDP is normally used with the <code>sendto()</code>, <code>sendmsg()</code>, <code>recvfrom()</code>, and <code>recvmsg()</code> calls (see <code>send(3SOCKET)</code> and <code>recv(3SOCKET)</code>). If the <code>connect(3SOCKET)</code> call is used to fix the destination for future packets, then the <code>recv(3SOCKET)</code> or <code>read(2)</code> and <code>send(3SOCKET)</code> or <code>write(2)</code> calls may be used.</p> <p>UDP address formats are identical to those used by the Transmission Control Protocol (“TCP”). Like TCP, UDP uses a port number along with an IP or IPv6 address to identify the endpoint of communication. The UDP port number space is separate from the TCP port number space, that is, a UDP port may not be “connected” to a TCP port. The <code>bind(3SOCKET)</code> call can be used to set the local address and port number of a UDP socket. The local IP or IPv6 address may be left unspecified in the <code>bind()</code> call by using the special value <code>INADDR_ANY</code> for IP, or the unspecified address (all zeroes) for IPv6. If the <code>bind()</code> call is not done, a local IP or IPv6 address and port number will be assigned to the endpoint when the first packet is sent. Broadcast packets may be sent, assuming the underlying network supports this, by using a reserved “broadcast address.” This address is network interface dependent. Broadcasts may only be sent by the privileged user.</p> <p>IPv6 does not support broadcast addresses; their function is supported by IPv6 multicast addresses.</p> <p>Options at the IP level may be used with UDP; see <a href="#">ip(7P)</a> or <a href="#">ip6(7p)</a>.</p> <p>There are a variety of ways that a UDP packet can be lost or corrupted, including a failure of the underlying communication mechanism. UDP implements a checksum over the data portion of the packet. If the checksum of a received packet is in error, the packet will be dropped with no indication given to the user. A queue of received packets is provided for each UDP socket. This queue has a limited capacity. Arriving datagrams which will not fit within its <i>high-water</i> capacity are silently discarded.</p> <p>UDP processes Internet Control Message Protocol (“ICMP”) and Internet Control Message Protocol Version 6 (“ICMP6”) error messages received in response to UDP packets it has sent. See <a href="#">icmp(7P)</a> and <a href="#">icmp6(7p)</a>.</p> |

## udp(7P)

ICMP “source quench” messages are ignored. ICMP “destination unreachable,” “time exceeded” and “parameter problem” messages disconnect the socket from its peer so that subsequent attempts to send packets using that socket will return an error. UDP will not guarantee that packets are delivered in the order they were sent. As well, duplicate packets may be generated in the communication process.

ICMP6 “destination unreachable” packets are ignored unless the enclosed code indicates that the port is not in use on the target host, in which case, the application is notified. ICMP6 “parameter problem” notifications are similarly passed upstream. All other ICMP6 messages are ignored.

**SEE ALSO** [read\(2\)](#), [write\(2\)](#), [bind\(3SOCKET\)](#), [connect\(3SOCKET\)](#), [recv\(3SOCKET\)](#), [send\(3SOCKET\)](#), [icmp\(7P\)](#), [icmp6\(7P\)](#), [inet\(7P\)](#), [inet6\(7P\)](#), [ip\(7P\)](#), [ip6\(7P\)](#), [tcp\(7P\)](#)

Postel, Jon, *RFC 768, User Datagram Protocol*, Network Information Center, SRI International, Menlo Park, Calif., August 1980

**DIAGNOSTICS** A socket operation may fail if:

|               |                                                                                                                                                                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EISCONN       | A <code>connect()</code> operation was attempted on a socket on which a <code>connect()</code> operation had already been performed, and the socket could not be successfully disconnected before making the new connection.                                                                 |
| EISCONN       | A <code>sendto()</code> or <code>sendmsg()</code> operation specifying an address to which the message should be sent was attempted on a socket on which a <code>connect()</code> operation had already been performed.                                                                      |
| ENOTCONN      | A <code>send()</code> or <code>write()</code> operation, or a <code>sendto()</code> or <code>sendmsg()</code> operation not specifying an address to which the message should be sent, was attempted on a socket on which a <code>connect()</code> operation had not already been performed. |
| EADDRINUSE    | A <code>bind()</code> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.                                                                                                                                                    |
| EADDRNOTAVAIL | A <code>bind()</code> operation was attempted on a socket with a network address for which no network interface exists.                                                                                                                                                                      |
| EINVAL        | A <code>sendmsg()</code> operation with a non-NULL <code>msg_accrights</code> was attempted.                                                                                                                                                                                                 |
| EACCES        | A <code>bind()</code> operation was attempted with a “reserved” port number and the effective user ID of the process was not the privileged user.                                                                                                                                            |
| ENOBUFS       | The system ran out of memory for internal data structures.                                                                                                                                                                                                                                   |

|                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|--------------------------------------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|-------------------------------------------------------------------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                                                                  | ufs – UFS file system                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <b>SYNOPSIS</b>                                                              | <pre>#include &lt;sys/param.h&gt; #include &lt;sys/types.h&gt; #include &lt;sys/fs/ufs_fs.h&gt; #include &lt;sys/fs/ufs_inode.h&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <b>DESCRIPTION</b>                                                           | <p>UFS is the default disk-based file system for the Solaris environment. The UFS file system is hierarchical, starting with its root directory (/) and continuing downward through a number of directories. The root of a UFS file system is <code>inode 2</code>. A UFS file system's root contents replace the contents of the directory upon which it is mounted.</p> <p>Subsequent sections of this manpage provide details of the UFS file systems.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <b>State Flags<br/>(<code>fs_state</code> and<br/><code>fs_clean</code>)</b> | <p>UFS uses state flags to identify the state of the file system. <code>fs_state</code> is <code>FSOKAY</code> - <code>fs_time</code>. <code>fs_time</code> is the timestamp that indicates when the last system write occurred. <code>fs_state</code> is updated whenever <code>fs_clean</code> changes. Some <code>fs_clean</code> values are:</p> <table border="0"> <tr> <td style="vertical-align: top;"><code>FSCLEAN</code></td> <td>Indicates an undamaged, cleanly unmounted file system.</td> </tr> <tr> <td style="vertical-align: top;"><code>FSACTIVE</code></td> <td>Indicates a mounted file system that has modified data in memory. A mounted file system with this state flag indicates that user data or metadata would be lost if power to the system is interrupted.</td> </tr> <tr> <td style="vertical-align: top;"><code>FSSTABLE</code></td> <td>Indicates an idle mounted file system. A mounted file system with this state flag indicates that neither user data nor metadata would be lost if power to the system is interrupted.</td> </tr> <tr> <td style="vertical-align: top;"><code>FSBAD</code></td> <td>Indicates that this file system contains inconsistent file system data.</td> </tr> <tr> <td style="vertical-align: top;"><code>FSLOG</code></td> <td>Indicates that the file system has logging enabled. A file system with this flag set is either mounted or unmounted. If a file system has logging enabled, the only flags that it can have are <code>FSLOG</code> or <code>FSBAD</code>. A non-logging file system can have <code>FSACTIVE</code>, <code>FSSTABLE</code>, or <code>FSCLEAN</code>.</td> </tr> </table> <p>It is not necessary to run the <code>fsck</code> command on unmounted file systems with a state of <code>FSCLEAN</code>, <code>FSSTABLE</code>, or <code>FSLOG</code>. <code>mount(2)</code> returns <code>ENOSPC</code> if an attempt is made to mount a UFS file system with a state of <code>FSACTIVE</code> for read/write access.</p> | <code>FSCLEAN</code> | Indicates an undamaged, cleanly unmounted file system. | <code>FSACTIVE</code> | Indicates a mounted file system that has modified data in memory. A mounted file system with this state flag indicates that user data or metadata would be lost if power to the system is interrupted. | <code>FSSTABLE</code> | Indicates an idle mounted file system. A mounted file system with this state flag indicates that neither user data nor metadata would be lost if power to the system is interrupted. | <code>FSBAD</code> | Indicates that this file system contains inconsistent file system data. | <code>FSLOG</code> | Indicates that the file system has logging enabled. A file system with this flag set is either mounted or unmounted. If a file system has logging enabled, the only flags that it can have are <code>FSLOG</code> or <code>FSBAD</code> . A non-logging file system can have <code>FSACTIVE</code> , <code>FSSTABLE</code> , or <code>FSCLEAN</code> . |
| <code>FSCLEAN</code>                                                         | Indicates an undamaged, cleanly unmounted file system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <code>FSACTIVE</code>                                                        | Indicates a mounted file system that has modified data in memory. A mounted file system with this state flag indicates that user data or metadata would be lost if power to the system is interrupted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <code>FSSTABLE</code>                                                        | Indicates an idle mounted file system. A mounted file system with this state flag indicates that neither user data nor metadata would be lost if power to the system is interrupted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <code>FSBAD</code>                                                           | Indicates that this file system contains inconsistent file system data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |
| <code>FSLOG</code>                                                           | Indicates that the file system has logging enabled. A file system with this flag set is either mounted or unmounted. If a file system has logging enabled, the only flags that it can have are <code>FSLOG</code> or <code>FSBAD</code> . A non-logging file system can have <code>FSACTIVE</code> , <code>FSSTABLE</code> , or <code>FSCLEAN</code> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                      |                                                        |                       |                                                                                                                                                                                                        |                       |                                                                                                                                                                                      |                    |                                                                         |                    |                                                                                                                                                                                                                                                                                                                                                        |

ufs(7FS)

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                         | <p>As an additional safeguard, <code>fs_clean</code> should be trusted only if <code>fs_state</code> contains a value equal to <code>FSOKAY - fs_time</code>, where <code>FSOKAY</code> is a constant integer defined in the <code>/usr/include/sys/fs/ufs_fs.h</code> file. Otherwise, <code>fs_clean</code> is treated as though it contains the value of <code>FSACTIVE</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Extended Fundamental Types (EFT)</b> | <p>Extended Fundamental Types (EFT) provide 32-bit user ID (UID), group ID (GID), and device numbers.</p> <p>If a UID or GID contains an extended value, the short variable (<code>ic_suid</code>, <code>ic_sgid</code>) contains the value 65535 and the corresponding UID or GID is in <code>ic_uid</code> or <code>ic_gid</code>. Because numbers for block and character devices are stored in the first direct block pointer of the inode (<code>ic_db[0]</code>) and the disk block addresses are already 32 bit values, no special encoding exists for device numbers (unlike UID or GID fields).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Multiterabyte File System</b>        | <p>A multiterabyte file system enables creation of a UFS file system up to approximately 16 terabytes of usable space, minus approximately one percent overhead. A sparse file can have a logical size of one terabyte. However, the actual amount of data that can be stored in a file is approximately one percent less than one terabyte because of file system overhead.</p> <p>On-disk format changes for a multiterabyte UFS file system include:</p> <ul style="list-style-type: none"><li>■ The magic number in the superblock changes from <code>FS_MAGIC</code> to <code>MTB_UFS_MAGIC</code>. For more information, see the <code>/usr/include/sys/fs/ufs_fs</code> file.</li><li>■ The <code>fs_logbno</code> unit is a sector for UFS that is less than 1 terabyte in size and fragments for a multiterabyte UFS file system.</li></ul>                                                                                                                                                                                                                         |
| <b>UFS Logging</b>                      | <p>UFS logging bundles the multiple metadata changes that comprise a complete UFS operation into a transaction. Sets of transactions are recorded in an on-disk log and are applied to the actual UFS file system's metadata.</p> <p>UFS logging provides two advantages:</p> <ol style="list-style-type: none"><li>1. A file system that is consistent with the transaction log eliminates the need to run <code>fsck</code> after a system crash or an unclean shutdown.</li><li>2. UFS logging often provides a significant performance improvement. This is because a file system with logging enabled converts multiple updates to the same data into single updates, thereby reducing the number of overhead disk operations.</li></ol> <p>The UFS log is allocated from free blocks on the file system, and is sized at approximately 1 Mbyte per 1 Gbyte of file system, up to a maximum of 64 Mbytes. The log is continually flushed as it fills up. The log is also flushed when the file system is unmounted or as a result of a <code>lockfs</code> command.</p> |
| <b>Mounting UFS File Systems</b>        | <p>You can mount a UFS file system in various ways using syntax similar to the following:</p> <ol style="list-style-type: none"><li>1. Use <code>mount</code> from the command line:</li></ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

```
mount -F ufs /dev/dsk/c0t0d0s7 /export/home
```

2. Include an entry in the `/etc/vfstab` file to mount the file system at boot time:

```
/dev/dsk/c0t0d0s7 /dev/rdisk/c0t0d0s7 /export/home ufs 2 yes -
```

For more information on mounting UFS file systems, see `mount_ufs(1M)`.

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Interface Stability | Unstable        |

**SEE ALSO** `fscck_ufs(1M)`, `mkfs_ufs(1M)`, `newfs(1M)`, `ufsdump(1M)`, `ufsrestore(1M)`, `tunefs(1M)`, `mount(2)`, `attributes(5)`

*Writing Device Drivers*

**NOTES** For information about internal UFS structures, see `newfs(1M)` and `mkfs_ufs(1M)`. For information about the `ufsdump` and `ufsrestore` commands, see `ufsdump(1M)`, `ufsrestore(1M)`, and `/usr/include/protocols/dumprestore.h`.

## ugen(7D)

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | ugen – USB generic driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>SYNOPSIS</b>    | Node Name@unit-address<br><br>#include <sys/usb/clients/ugen/usb_ugen.h>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>DESCRIPTION</b> | <p>ugen is a generic USBA (Solaris USB Architecture) compliant client character driver that presents USB devices to applications through a standard <code>open(2)</code>, <code>close(2)</code>, <code>read(2)</code>, <code>write(2)</code>, <code>aioread(3AIO)</code>, <code>aiowrite(3AIO)</code> Unix interface. Uninterpreted raw data are transferred to and from the device via file descriptors created for each USB endpoint. Status is obtained by reading file descriptors created for endpoint and full device status.</p> <p>ugen supports control, bulk, and interrupt (in and out) transfers. Isochronous transfers are not supported.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>BINDING</b>     | <p>ugen can bind to a device with one or more interfaces in its entirety, or to a single interface of that device. The binding type depends on information that is passed to <code>add_drv(1M)</code> or <code>update_drv(1M)</code>.</p> <p>An <code>add_drv(1M)</code> command binds ugen to a list of device types it is to control. <code>update_drv(1M)</code> adds an additional device type to the list of device types being managed by the driver.</p> <p>Names used to bind drivers can be found in <code>/var/adm/messages</code>. When a device is onlined after hot insertion, and no driver is found, there will be an entry containing:</p> <pre>USB 2.0 device (usb&lt;vid&gt;,&lt;pid&gt;)...</pre> <p>where <code>vid</code> is the USB vendor identifier and <code>pid</code> is the product identifier supplied by the device descriptor <code>usb_dev_descr(9S)</code>.</p> <p>When using ugen for the first time, you must add the driver utilizing <code>add_drv(1M)</code>, using a command of the following form:</p> <pre>add_drv -n -m '* &lt;device perms&gt; &lt;owner&gt; &lt;group&gt;' -i 'usb1234,5678' ugen</pre> <p>If the command fails with:</p> <pre>(ugen) already in use as a driver or alias.</pre> <p>...add the device using <code>update_drv(1M)</code>:</p> <pre>update_drv -n -m '* &lt;device perms&gt; &lt;owner&gt; &lt;group&gt;' -i 'usb1234,5678' ugen</pre> <p>This binds ugen to the entire device.</p> <p>If ugen only binds to one interface of the device, use the following <code>driver_alias</code> instead of <code>usb&lt;vid&gt;,&lt;pid&gt;</code>:</p> <pre>usbif&lt;vid&gt;,&lt;pid&gt;.config&lt;cfg value&gt;.&lt;interface number&gt;</pre> <p>where <code>cfg value</code> is the value of <code>bConfigurationValue</code> in the configuration descriptor (<code>usb_cfg_descr(9S)</code>). For example "usbif1234,4567.config1.0."</p> |



Note that you can use `update_drv` to also remove bindings. Please see `update_drv(1M)` for more information.

After a successful `add_drv` or `update_drv`, remove the device and reinsert. Check with the `prtconf(1M)` `-D` option to determine if `ugen` is successfully bound to the device and the nodes created in `/dev/usb/<vid>.<pid>` (see below).

An example showing how to bind a child device representing interface 0 of configuration 1 of a composite device follows:

```
update_drv -a -m '* 0666 root sys'
-i '"usbif472,b0b0.config1.0"' ugen
```

Note that you can completely uninstall the `ugen` driver and delete it from the system by doing:

```
pkgrm SUNWugen
```

Any `pkgadd` of `SUNWugen` after the `pkgrm` reactivates any pre-existing `ugen` driver device-bindings.

Any pre-existing `ugen` driver device-bindings are preserved across operating system upgrades.

#### LOGICAL DEVICE NAME FORMAT

For each device or child device it manages, `ugen` creates one logical device name for device-wide status and one logical device name for endpoint 0. `ugen` also creates logical device names for all other endpoints within the device node's binding scope (interface or device), plus logical device names for their status.

If separate `ugen` instances control different interfaces of the same device, the device-wide status and endpoint logical device names created for each instance will share access to the same source or endpoint pipes. For example, a device with two interfaces, each operated by their own `ugen` instance, will show `endpoint0` as `if0cntrl0` to the first interface, and will show it as `if1cntrl0` to the second interface. Both of these logical device names share `endpoint0`. Likewise for the same device, `ugen` makes the device-wide status available as `if0devstat` to the first interface and as `if1devstat` to the second interface. `if0devstat` and `if1devstat` both return the same data.

Any `ugen` logical device name can be held open by only one user at a time, regardless of whether the `O_EXCL` flag passed to `open(2)`. When a single pipe or data source is shared by multiple logical device names, such as `if[0,1]cntrl0` or `if[0,1]devstat` above, more than one logical device name sharing the pipe or data source can be open at a time. However, only one user may access the shared pipe or data source at a time, regardless of the logical device name used for access.

When `ugen` is bound to an entire device, the following logical device names are created (each on a single line). *N* represents the instance number of the device type.

```
Endpoint 0 (default endpoint):
/dev/usb/<vid>.<pid>/<N>/cntrl0
```

## ugen(7D)

```
/dev/usb/<vid>.<pid>/<N>/cntrl0stat
```

For example:

```
/dev/usb/472.b0b0/0/cntrl0
/dev/usb/472.b0b0/0/cntrl0stat
```

Configuration index 1, Endpoints > 0, alternate 0:

```
/dev/usb/<vid>.<pid>/<N>/if<interface#>
 <in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/if<interface#>
 <in|out|cntrl><endpoint#>stat
```

For example:

```
/dev/usb/472.b0b0/0/if0in1
/dev/usb/472.b0b0/0/if0in1stat
```

Configuration index 1, Endpoints > 0, alternate > 0:

```
/dev/usb/<vid>.<pid>/<N>/if<interface#>.
 <alternate><in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/if<interface#>.
 <alternate><in|out|cntrl><endpoint#>stat
```

For example:

```
/dev/usb/472.b0b0/0/if0.1in3
/dev/usb/472.b0b0/0/if0.1in3stat
```

Configuration index> 1, Endpoints > 0, alternate 0:

```
/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>
 <in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>
 <in|out|cntrl><endpoint#>stat
```

For example:

```
/dev/usb/472.b0b0/0/cfg2if0in1
/dev/usb/472.b0b0/0/cfg2if0in1stat
```

Note that the configuration value from the configuration descriptor indexed by the configuration index is used in the node name and not the configuration index itself.

Configuration index> 1, Endpoints > 0, alternate > 0:

```
/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>.
 <alternate><in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>.
 <alternate><in|out|cntrl><endpoint#>stat
```

For example:

```
/dev/usb/472.b0b0/0/cfg2if0.1in1
/dev/usb/472.b0b0/0/cfg2if0.1in1stat
```

Device status:

```
/dev/usb/<vid>.<pid>/<N>/devstat
```

For example:

```
/dev/usb/472.b0b0/0/devstat
```

When `ugen` is bound to a single device interface, the following logical device nodes are created:

Endpoint 0 (default endpoint):

```
/dev/usb/<vid>.<pid>/<N>/if<interface#>cntrl0
/dev/usb/<vid>.<pid>/<N>/if<interface#>cntrl0stat
```

For example:

```
/dev/usb/472.b0b0/0/if0cntrl0
/dev/usb/472.b0b0/0/if0cntrl0stat
```

Device status:

```
/dev/usb/<vid>.<pid>/<N>/if<interface#>devstat
```

For example:

```
/dev/usb/472.b0b0/0/if0devstat
```

The format for all other logical device names is identical to the format used when `ugen` is bound to the entire device.

Opening the endpoint of a different configuration or different alternate interface will cause an implicit change of configuration or a switch to an alternate interface. A configuration change is prohibited when any non-zero endpoint device nodes are open. An alternate interface switch is prohibited if any endpoint in the same interface is open.

## HOT-PLUGGING

A device may be hot-removed at any time. Following hot-removal, the device status changes to `USB_DEV_STAT_DISCONNECTED`, the status of open endpoints change to `USB_LC_STAT_DISCONNECTED` upon their access, and all subsequent transfer requests fail. Endpoints are reactivated by first reinserting the device and then closing and reopening all endpoints that were open when the device was disconnected.

## CPR (CHECKPOINT/RESUME)

CPR (Checkpoint/Resume) may be initiated at any time and is treated similarly to a hot-removal. Upon successful suspend and resume, all subsequent transfer requests fail as an indication to the application to reinitialize. Applications should close and reopen all endpoints to reinstate them. All endpoint and device status on Resume (before close and reopen) is `USB_LC_STAT_SUSPENDED`. A system suspend will fail while `ugen` is performing a transfer.

## DEVICE POWER MANAGEMENT

Devices which support remote wakeup can be power managed when they have no open logical device nodes. When an application opens the first logical device node of a device, that application should assume that a reinitialization of device state is required.

ugen(7D)

## DEVICE STATUS MANAGEMENT

Applications can monitor device status changes by reading the device status from the device status logical name. When opened without `O_NONBLOCK` and `O_NDELAY`, all reads from that file descriptor (with the exception of the the intial read that follows the open) block until a device status change occurs. Calls to read will always return immediately if opened with `O_NONBLOCK` or `O_NDELAY`. Nonblocking calls to read which have no data to return, return no error and zero bytes read.

Device statuses are:

`USB_DEV_STAT_ONLINE`

Device is available.

`USB_DEV_STAT_DISCONNECTED`

Device has been disconnected.

`USB_DEV_STAT_RESUMED`

Device has been resumed, however, endpoints which were open on suspend have not yet been closed and reopened.

`USB_DEV_STAT_UNAVAILABLE`

Device has been reconnected, however, endpoints which were open on disconnect have not yet been closed and reopened.

The following code reads the device status device logical name:

```
int fd;
int status;

if ((fd = open("/dev/usb/472.b0b0/0/devstat",
 O_RDONLY)) < 0) {
 /* handle error */
}

if (read(fd, &status, sizeof(status)) != sizeof(status)) {
 /* handle error */
}

switch (status) {
case USB_DEV_STAT_DISCONNECTED:
 printf ("Terminating as device has been disconnected.\n");
 exit (0);

case USB_DEV_STAT_RESUMED:
case USB_DEV_STAT_UNAVAILABLE:
 /*
 * Close and reopen endpoints to reestablish device access,
 * then reset device.
 */
 break;

case USB_DEV_STAT_ONLINE:
default:
 break;
}
```

**ENDPOINT  
STATUS  
MANAGEMENT**

Use `poll(2)` to block on several logical names simultaneously, including device status logical names. Poll indicates when reading a logical name would return data. See `poll(2)` for details. Calls to read may be done whether or not they follow calls to poll.

Each data endpoint has a corresponding status logical name. Use the status logical name to retrieve the state of the data endpoint, including detail on how its most recent transfer failed. Reads of the status file descriptors always return immediately. See the ERRORS section for more information on endpoint status values. All logical device name files created for returning status must be opened with `O_RDONLY`.

The following code illustrates reading the status file descriptor of an endpoint which just failed a data transfer in order to get more information on the failure.

```
int data_xfered, status;
int ep1_data_fd, ep1_stat_fd;
uchar_t request[8];

ep1_data_fd = open ("/dev/usb/472.b0b0/0/if0out1", O_WRONLY);

if (ep1_data_fd < 0) {
 /* Handle open error. */
}

ep1_stat_fd = open ("/dev/usb/472.b0b0/0/if0out1stat",
 O_RDONLY);
if (ep1_stat_fd < 0) {
 /* Handle open error. */
}

data_xfered = write(ep1_data_fd, request, sizeof (request));

/* An error occurred during the data transfer. */
if (data_xfered != sizeof (request)) {

 /* Read status file descriptor for details on failure. */
 if (read(ep1_stat_fd, (int *)&status, sizeof (status)) !=
 sizeof (status)) {
 status = USB_LC_STAT_UNSPECIFIED_ERR;
 }

 /* Take appropriate action. */
 switch (status) {
 case USB_LC_STAT_STALL:
 printf ("Endpoint stalled.\n");
 break;
 case ...
 ...
 }
}
}
```

**CONTROL  
TRANSFERS**

The control endpoint is typically used to set up the device and to query device status or configuration.

Applications requiring I/O on a control endpoint should open the corresponding logical device name and use regular UNIX I/O system calls. For example: `read(2)`, `write(2)`, `aioread(3AIO)` and `aiowrite(3AIO)`. `poll(2)` is not supported on control endpoints.

A control endpoint must be opened with `O_RDWR` since it is bidirectional. It cannot be opened with `O_NONBLOCK` or `O_NDELAY`.

For example:

```
fd = open("/dev/usb/472.b0b0/0/cntrl0", O_RDWR);
```

```
fdstat = open("/dev/usb/472.b0b0/0/cntrl0stat", O_RDONLY);
```

Control endpoints can be read and written. A read operation receives data *from* the device and a write operation sends data *to* the device.

To perform a control-IN transfer, perform a `write(2)` of USB setup data (see section 9.3 of the *USB 1.1* or *2.0* specifications) followed by a `read(2)` on the same control endpoint to fetch the desired data. For example:

```
void init_cntrl_req(
 uchar_t *req, uchar_t bmRequestType, uchar_t bRequest,
 ushort_t wValue, ushort_t wIndex, ushort_t wLength) {
 req[0] = bmRequestType;
 req[1] = bRequest;
 req[2] = 0xFF & wValue;
 req[3] = 0xFF & (wValue >> 8);
 req[4] = 0xFF & wIndex;
 req[5] = 0xFF & (wIndex >> 8);
 req[6] = 0xFF & wLength;
 req[7] = 0xFF & (wLength >> 8);
}

....

uchar_t dev_descr_req[8];
usb_dev_descr_t descr;

init_cntrl_req(dev_descr_req,
 USB_DEV_REQ_DEV_TO_HOST, USB_REQ_GET_DESCR,
 USB_DESCR_TYPE_SETUP_DEV, 0, sizeof (descr));

count = write(fd, dev_descr_req, sizeof (dev_descr_req));
if (count != sizeof (dev_descr_req)) {
 /* do some error recovery */
 ...
}

count = read(fd, &descr, sizeof (descr));
if (count != sizeof (descr)) {
 /* do some error recovery */
}
}
```

The application can issue any number of reads to read data received on a control endpoint. `ugen` successfully completes all reads, returning the number of bytes transferred. Zero is returned when there is no data to transfer.

If the read/write fails and returns `-1`, you can access the endpoint's status device logical name for precise error information:

```
int status;

count = read(fdstat, &status, sizeof (status));
if (count == sizeof (status)) {
 switch (status) {
 case USB_LC_STAT_SUSPENDED:
 case USB_LC_STAT_DISCONNECTED:
 /* close all endpoints */
 ...
 break;
 default:
 ...
 break;
 }
}
```

Refer to the **ERRORS** section for all possible error values.

To perform a control-OUT transfer, send in a single transfer, the USB setup data followed by any accompanying data bytes.

```
/* 1st 8 bytes of wbuf are setup. */
init_cntrl_req(wbuf,);

/* Data bytes begin at byte 8 of wbuf. */
bcopy(data, &wuf[8], sizeof (data));

/* Send it all in a single transfer. */
count = write(fd, wbuf, sizeof (wbuf));
```

A `write(2)` returns the number of bytes (both setup and data) actually transferred, (whether or not the write is completely successful), provided that some data is actually transferred. When no data is transferred, `write(2)` returns `-1`. Applications can read the corresponding endpoint status to retrieve detailed error information. Note that it is an error to specify a size different than:

(number of data bytes + number of setup bytes).

Here is a more extensive example which gets all descriptors of a device configuration. For sake of brevity, uninteresting parts are omitted.

```
#include <sys/usb/usba.h>
#include <sys/usb/clients/ugen/usb_ugen.h>

uchar_t *config_cloud;
uchar_t *curr_descr;

uchar_t *bytes;
```

```

int curr_descr_len;
int curr_descr_type;

usb_cfg_descr_t cfg_descr;
usb_if_descr_t if_descr;
usb_ep_descr_t ep_descr;

/* See 9.13 of USB 2.0 spec for ordering. */
static char *pipetypes[] = {
 "Control", "Isochronous", "Bulk", "Interrupt"
};

/*
 * Setup to send a request to read just the config descriptor. The
 * size of the whole cloud, containing all cfg, interface, endpoint,
 * class and vendor-specific descriptors, will be returned as part of
 * the config descriptor.
 */
init_ctrl_req(&setup_data, USB_DEV_REQ_DEV_TO_HOST, USB_REQ_GET_DESCR,
 USB_DESCR_TYPE_SETUP_CFG, 0, USB_CFG_DESCR_SIZE);

/*
 * Write setup data. USB device will prepare to return the whole
 * config cloud as a response to this. We will read this separately.
 */
count = write(ctrl_fd, &setup_data, sizeof (setup_data));
if (count != sizeof (setup_data)) {
 /* Error recovery. */
} else {
 count = read(ctrl_fd, &cfg_descr, USB_CFG_DESCR_SIZE);
 if (count != USB_CFG_DESCR_SIZE) {
 /* Error recovery. */
 }
}

/* USB data is little endian. */
bytes = (uchar_t *)&cfg_descr.wTotalLength;
totalLength = bytes[0] + (bytes[1] << 8);

/*
 * The size of the whole cloud is in the bLength field. Set up
 * to read this amount of data, to get the whole cloud.
 */
config_cloud = malloc(totalLength);

init_ctrl_req(&setup_data, USB_DEV_REQ_DEV_TO_HOST, USB_REQ_GET_DESCR,
 USB_DESCR_TYPE_SETUP_CFG, 0, totalLength);

count = write(ctrl_fd, &setup_data, sizeof (setup_data));
if (count != sizeof (setup_data)) {
 /* Error recovery. */
} else {
 count = read(ctrl_fd, config_cloud, totalLength);
 if (count != totalLength) {
 /* Error recovery. */
 }
}
}

```



```

/* Got the data. Now loop, dumping out the descriptors found. */

curr_descr = config_ccloud;
offset = 0;
while (offset < totalLength) {

 /* All descr have length and type at offset 0 and 1 */
 curr_descr_len = curr_descr[0];
 curr_descr_type = curr_descr[1];

 switch (curr_descr_type) {
 case USB_DESCRIPTOR_TYPE_CFG:

 /*
 * Copy data into separate structure, needed for
 * proper alignment of all non char fields. Note:
 * non-char fields of all descriptors begin on aligned
 * boundaries. The issue is that some structures may
 * be adjacent to others which have an odd-numbered
 * byte size, and may thus start on an odd-numbered
 * boundary. */
 bcopy(curr_descr, &cfg_descr, curr_descr_len);

 /* Remember to read any words in endian-neutral way. */

 (void) printf("\nConfig %d found.\n",
 cfg_descr.bConfigurationValue);
 break;

 case USB_DESCRIPTOR_TYPE_IF:
 bcopy(curr_descr, &if_descr, curr_descr_len);
 (void) printf("\n\tInterface %d, Alt %d found.\n",
 if_descr.bInterfaceNumber,
 if_descr.bAlternateSetting);
 break;

 case USB_DESCRIPTOR_TYPE_EP:
 bcopy(curr_descr, &ep_descr, curr_descr_len);
 (void) printf("\n\t\tEndpoint %d (%s-%s) found.\n",
 (ep_descr.bEndpointAddress & USB_EP_NUM_MASK),
 (pipetypes[
 ep_descr.bmAttributes & USB_EP_ATTR_MASK]),
 ((ep_descr.bEndpointAddress &
 USB_EP_DIR_IN) ? "IN" : "OUT"));
 break;

 default:
 (void) printf(
 "\n\t\t\tOther descriptor found. Type:%d\n",
 curr_descr_type);
 break;
 }

 offset += curr_descr_len;
 curr_descr = &config_ccloud[offset];
}

```

ugen(7D)

## INTERRUPT-IN TRANSFERS

Applications requiring data from an interrupt-IN endpoint should open the corresponding logical device name and use `read(2)`, `aioread(3AIO)` and `poll(2)` system calls.

An interrupt-IN endpoint must be opened with `O_RDONLY`. It can also be opened using `O_NONBLOCK` or `O_NDELAY` if desired.

```
fd = open("/dev/usb/472.b0b0/0/if0in1", O_RDONLY);
fdstat = open("/dev/usb/472.b0b0/0/if0in1stat", O_RDONLY);
```

ugen starts polling interrupt-IN endpoints immediately upon opening them and stops polling them upon closure. (Polling refers to interrogation of the device by the driver and should not be confused with `poll(2)`, which is an interrogation of the driver by the application.)

A `read(2)` of an endpoint opened with the `O_NONBLOCK` or `O_NDELAY` flags set will not block when there is insufficient data available to satisfy the request. The `read` simply returns what it can without signifying any error.

Applications should continuously check for and consume interrupt data. ugen enables buffering of up to one second of incoming data. In case of buffer overflow, ugen stops polling the interrupt-IN endpoint until the application consumes all the data. In this case, a `read(2)` of an empty buffer returns `-1`, sets the endpoint status to `USB_LC_STAT_INTR_BUF_FULL` (to indicate that the buffer had been full and polling had been stopped) and causes ugen to start polling the endpoint again. To retrieve the status, the application can open and read the corresponding endpoint's status device logical name.

```
for (;;) {
 count = read(fd, buf, sizeof(buf));
 if (count == -1) {
 int cnt, status;

 cnt = read(fdstat, &status, sizeof(status));
 if (cnt == -1) {
 /* more error recovery here */
 } else {
 switch (status) {
 case USB_LC_STAT_INTR_BUF_FULL:
 ...
 break;
 default:
 ...
 break;
 }
 }
 }
 /* process the data */

}
```

ugen will never drop data. However, the device may drop data if the application cannot read it at the rate that it is produced.

Applications requiring unbuffered data from an interrupt-IN endpoint should open the associated status endpoint with `O_RDWR` before opening the associated interrupt-IN endpoint and write a control byte with `USB_EP_INTR_ONE_XFER` set. All other bits are reserved and should be 0.

"One transfer" mode will persist until disabled explicitly after the associated interrupt-IN endpoint has been closed by writing a control byte with `USB_EP_INTR_ONE_XFER` cleared.

"One transfer" mode is implicitly disabled when the status/control endpoint is closed.

Attempts to change the "one transfer" mode while the endpoint is open will result in `EINVAL`.

An application can open multiple interrupt-IN endpoints and can call `poll(2)` to monitor the availability of new data. (Note: `poll` works with interrupt-IN data endpoints, not their status endpoints.)

```

struct pollfd pfd[2];

bzero(pfd, sizeof (pfd));
pfd[0].fd = fd1; /* fd1 is one interrupt-IN endpoint. */
pfd[0].events = POLLIN;
pfd[1].fd = fd2; /* fd2 is another interrupt-IN endpoint. */
pfd[1].events = POLLIN;

for (;;) {
 poll(pfd, 2, -1);

 if (pfd[0].revents & POLLIN) {
 count = read(fd1, buf, sizeof (buf));

 }
 if (pfd[1].revents & POLLIN) {
 count = read(fd2, buf, sizeof (buf));

 }
}

```

You can monitor the device status endpoint via `poll(2)` concurrently with the multiple interrupt-IN endpoints. Simply add another `pollfd` element to the `pfd` array in the previous code example, and initialize the new element's `fd` field with the file descriptor of the device status endpoint (opened without `O_NONBLOCK` or `O_NDELAY`). Set the new element's event field to `POLLIN` like the other elements. Note that only interrupt-IN endpoints and the device status endpoint can be monitored using `poll(2)`.

## INTERRUPT-OUT TRANSFERS

Applications requiring output on an interrupt-OUT endpoint can open the corresponding logical device name and perform regular UNIX I/O system calls such as `write(2)` and `aiowrite(3AIO)`.

An interrupt-OUT endpoint must be opened with `O_WRONLY`.

```
fd = open("/dev/usb/472.b0b0/0/if0out3", O_WRONLY);
```

ugen(7D)

```
fdstat = open("/dev/usb/472.b0b0/0/if0out3stat", O_RDONLY);
```

Data can be written to an interrupt-OUT endpoint as follows:

```
count = write(fd, buf, sizeof (buf));
if (count == -1) {
 /* error recovery */
}
```

## **BULK TRANSFERS**

Applications requiring I/O on a bulk endpoint can open the corresponding logical device name and perform regular UNIX I/O system calls. For example: `read(2)`, `write(2)`, `aioread(3AIO)` and `aiowrite(3AIO)`. `poll(2)` is not supported on bulk endpoints.

A bulk endpoint must be opened with `O_RDONLY` or `O_WRONLY` and cannot be opened with `O_NONBLOCK` or `O_NDELAY`:

```
fd = open("/dev/usb/472.b0b0/0/if0in2", O_RDONLY);
fdstat = open("/dev/usb/472.b0b0/0/if0in2stat", O_RDONLY);
```

Data can be read from a bulk-IN endpoint as follows:

```
count = read(fd, buf, sizeof (buf));
if (count == -1) {
 /* error recovery */
}
```

Data can be written to a bulk-OUT endpoint as follows:

```
count = write(fd, buf, sizeof (buf));
if (count == -1) {
 /* error recovery */
}
```

## **ERRORS**

The following statuses are returned by endpoint status device logical names:

`USB_LC_STAT_NOERROR`

No error.

`USB_LC_STAT_CRC`

CRC error detected.

`USB_LC_STAT_BITSTUFFING`

Bit stuffing error.

`USB_LC_STAT_DATA_TOGGLE_MM`

Data toggle did not match.

`USB_LC_STAT_STALL`

Endpoint returned stall.

`USB_LC_STAT_DEV_NOT_RESP`

Device not responding.

USB\_LC\_STAT\_UNEXP\_PID  
 Unexpected Packet Identifier (PID).

USB\_LC\_STAT\_PID\_CHECKFAILURE  
 Check bits on PID failed.

USB\_LC\_STAT\_DATA\_OVERRUN  
 Data overrun.

USB\_LC\_STAT\_DATA\_UNDERRUN  
 Data underrun.

USB\_LC\_STAT\_BUFFER\_OVERRUN  
 Buffer overrun.

USB\_LC\_STAT\_BUFFER\_UNDERRUN  
 Buffer underrun.

USB\_LC\_STAT\_TIMEOUT  
 Command timed out.

USB\_LC\_STAT\_NOT\_ACCESSED  
 Not accessed by the hardware.

USB\_LC\_STAT\_UNSPECIFIED\_ERR  
 Unspecified USB A or HCD error.

USB\_LC\_STAT\_NO\_BANDWIDTH  
 No bandwidth available.

USB\_LC\_STAT\_HW\_ERR  
 Host Controller h/w error.

USB\_LC\_STAT\_SUSPENDED  
 Device was suspended.

USB\_LC\_STAT\_DISCONNECTED  
 Device was disconnected.

USB\_LC\_STAT\_INTR\_BUF\_FULL  
 Polling was stopped as the interrupt-IN data buffer was full. Buffer is now empty and polling has been resumed.

USB\_LC\_STAT\_INTERRUPTED  
 Request was interrupted.

USB\_LC\_STAT\_NO\_RESOURCES  
 No resources available for request.

USB\_LC\_STAT\_INTR\_POLLING\_FAILED  
 Failed to restart polling.

The following system call `errno` values are returned:

|                     |                                                                                                      |
|---------------------|------------------------------------------------------------------------------------------------------|
| <code>EINVAL</code> | An attempt was made to enable or disable "one transfer" mode while the associated endpoint was open. |
|---------------------|------------------------------------------------------------------------------------------------------|

## ugen(7D)

|         |                                                                                                          |
|---------|----------------------------------------------------------------------------------------------------------|
| EBUSY   | The endpoint has been opened and another open is attempted.                                              |
| EACCES  | An endpoint open was attempted with incorrect flags.                                                     |
| ENOTSUP | Operation not supported.                                                                                 |
| ENXIO   | Device associated with the file descriptor does not exist.                                               |
| ENODEV  | Device has been hot-removed or a suspend/resume happened before this command.                            |
| EIO     | An I/O error occurred. Send a read on the endpoint status minor node to get the exact error information. |
| EINTR   | Interrupted system call.                                                                                 |
| ENOMEM  | No memory for the allocation of internal structures.                                                     |

### FILES

```
/kernel/drv/ugen 32 bit ELF kernel module (x86 platform only)
/kernel/drv/sparcv9/ugen 64 bit ELF kernel module

/dev/usb/<vid>.<pid>/<N>/cntrl0
/dev/usb/<vid>.<pid>/<N>/cntrl0stat

/dev/usb/<vid>.<pid>/<N>/if<interface#>
 <in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/if<interface#>
 <in|out|cntrl><endpoint#>stat

/dev/usb/<vid>.<pid>/<N>/if<interface#>.
 <alternate><in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/if<interface#>.
 <alternate><in|out|cntrl><endpoint#>stat

/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>
 <in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>
 <in|out|cntrl><endpoint#>stat

/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>.
 <alternate><in|out|cntrl><endpoint#>
/dev/usb/<vid>.<pid>/<N>/cfg<value>if<interface#>.
 <alternate><in|out|cntrl><endpoint#>stat

/dev/usb/<vid>.<pid>/<N>/devstat

/dev/usb/<vid>.<pid>/<N>/if<interface#>cntrl0
/dev/usb/<vid>.<pid>/<N>/if<interface#>cntrl0stat
```

where *N* is an integer representing the instance number of this type of device. (All logical device names for a single device share the same *N*.)

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Architecture   | PCI-based SPARC |
| Availability   | SUNWusb         |

**SEE ALSO** `usba(7D)`, `usb_dev_descr(9S)`

**DIAGNOSTICS** In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

Warning: <device path> (ugen<instance num>): Error Message...

Too many minor nodes.

Device has too many minor nodes. Not all are available.

Instance number too high (<number>).

Too many devices are using this driver.

Cannot access <device>. Please reconnect.

This device has been disconnected because a device other than the original one has been inserted. The driver informs you of this fact by displaying the name of the original device.

Device is not identical to the previous one on this port. Please disconnect and reconnect.

Same condition as described above; however in this case, the driver is unable to identify the original device with a name string.

**NOTES** Isochronous transfers are not supported.

ugen returns -1 for all commands and sets `errno` to `ENODEV` when device has been hot-removed or resumed from a suspend. The application must close and reopen all open minor nodes to reinstate successful communication.

uhci(7D)

**NAME** | uhci – host controller driver

**SYNOPSIS** | pcivid, pid@unit-address

**DESCRIPTION** | The uhci host controller driver is a USBA (Solaris USB Architecture) compliant nexus driver that supports the *Universal Host Controller Interface Specification 1.1*, an industry standard developed by Intel. The uhci driver supports all USB transfers, including interrupt, control, isochronous and bulk.

**FILES** | /kernel/drv/uhci            32-bit x86 ELF kernel module  
          | /kernel/drv/amd64/uhci    64-bit x86 ELF kernel module  
          | /kernel/drv/uhci.conf    driver configuration file

**ATTRIBUTES** | See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE        |
|----------------|------------------------|
| Architecture   | x86, PCI-based systems |
| Availability   | SUNWusb                |

**SEE ALSO** | attributes(5), ehci(7D), hubd(7D), ohci(7D), usba(7D)

          | *Writing Device Drivers*

          | *Universal Host Controller Interface Specification for USB 1.1*

          | *Universal Serial Bus Specification 2.0*

          | *System Administration Guide: Basic Administration*

          | <http://www.sun.com/io>

**DIAGNOSTICS** | All host controller errors are passed to the client drivers. Root errors are documented in hubd(7D).

          | In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

          | WARNING: <device path> <uhci><instance number>>: Error message...

          | No SOF interrupts have been received. This USB UHCI host controller is unusable. The USB hardware is not generating Start Of Frame interrupts. Please reboot the system. If this problem persists, contact your system vendor.



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | usba, usb – Solaris USB Architecture (USBA)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>DESCRIPTION</b> | <p>USB provides a low-cost means for attaching peripheral devices, including mass-storage devices, keyboards, mice, and printers, to a system. For complete information on USB, go to the USB website at <a href="http://www.usb.org">http://www.usb.org</a>.</p> <p>USBA supports 126 hot-pluggable USB devices per USB bus. The maximum data transfer rate is 1.5 Mbits (low speed USB 1.x) or 12 Mbits (full speed USB 1.x) or 480 Mbits (high speed USB 2.0) Mbits per second (Mbps).</p> <p>USBA adheres to the <i>Universal Serial Bus 2.0</i> specification and provides a transport layer abstraction to USB client drivers.</p> |
| <b>FILES</b>       | Listed below are drivers and modules which either utilize or are utilized by USBA. Drivers in <code>/kernel/drv</code> are 32 bit drivers (x86 only). Drivers in <code>/kernel/drv/sparcv9</code> or <code>/kernel/drv/amd64</code> are 64 bit drivers.                                                                                                                                                                                                                                                                                                                                                                                  |

| CLIENT DRIVER                                        | FUNCTION/DEVICE                  |
|------------------------------------------------------|----------------------------------|
| <code>/kernel/drv/[sparcv9 amd64]/hid</code>         | HID class                        |
| <code>/kernel/drv/[sparcv9 amd64]/hubd</code>        | hub class                        |
| <code>/kernel/drv/[sparcv9 amd64]/scsa2usb</code>    | mass storage class               |
| <code>/kernel/drv/[sparcv9 amd64]/usbprn</code>      | printer class                    |
| <code>/kernel/drv/[sparcv9 amd64]/usb_as</code>      | audio streaming class            |
| <code>/kernel/drv/[sparcv9 amd64]/usb_ac</code>      | audio control class              |
| <code>/kernel/drv/[sparcv9 amd64]/usb_mid</code>     | multi-interface device           |
| <code>/kernel/drv/[sparcv9 amd64]/usbser_edge</code> | Edgeport USB to serial port      |
| <code>/kernel/drv/[sparcv9 amd64]/ugen</code>        | generic USB driver               |
| <code>/kernel/drv/[sparcv9 amd64]/ohci</code>        | open host controller driver      |
| <code>/kernel/drv/[sparcv9 amd64]/uhci</code>        | universal host controller driver |
| <code>/kernel/drv/[sparcv9 amd64]/ehci</code>        | enhanced host controller driver  |

| CLIENT STREAMS MODULES                       | FUNCTION/DEVICE |
|----------------------------------------------|-----------------|
| <code>/kernel/strmod/[sparcv9/]usbkbm</code> | Keyboard        |
| <code>/kernel/strmod/[sparcv9/]usbms</code>  | Mouse           |
| <code>/kernel/strmod/[sparcv9/]usb_ah</code> | audio HID       |

usba(7D)

| HOST CONTROLLER INTERFACE DRIVERS | DEVICE        |
|-----------------------------------|---------------|
| /kernel/drv/[sparcv9/]ehci        | Enhanced HCI  |
| /kernel/drv/[sparcv9/]ohci        | Open HCI      |
| /kernel/drv/[sparcv9/]uhci        | Universal HCI |

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE   |
|----------------|-------------------|
| Architecture   | PCI-based systems |
| Availability   | SUNWusb           |

**SEE ALSO** `cfgadm_usb(1M)`, `attributes(5)`, `ehci(7D)`, `hid(7D)`, `hubd(7D)`, `ohci(7D)`, `scsa2usb(7D)`, `uhci(7D)`, `usb_ac(7D)`, `usb_as(7D)`, `usb_mid(7D)`, `usbprn(7D)`, `usbser_edge(7D)`, `ugen(7D)`

*Universal Serial Bus Specification 2.0.*

*System Administration Guide: Basic Administration*

<http://www.sun.com/io>

**NOTES** Booting from USB mass-storage devices is not supported.

**DIAGNOSTICS** The messages described below may appear on the system console as well as being logged. All messages are formatted in the following manner:

WARNING: Error message...

No driver found for device <device\_name> (interface <number> node name=<node\_name>)

The installed Solaris software does not contain a supported driver for this hardware. <number> is the interface number. <name> is either the device path name or the device name.

Draining callbacks timed out!

An internal error occurred. Please reboot your system. If this problem persists, contact your system vendor.

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                             | usb_ac – USB audio control driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>SYNOPSIS</b>                         | sound-control@unit-address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>DESCRIPTION</b>                      | <p>The usb_ac driver is a USBA (Solaris USB Architecture) compliant client driver that supports the <i>USB Audio Class 1.0</i> specification.</p> <p>The audio control driver is a USB class driver and offers functionality similar to the audiocs (sun4u) and audiotcs (Sun Blade 100) drivers which use the Solaris audio mixer framework (mixer(7I)). Unlike the audiocs and audiotcs drivers, the USB audio device may have play-only or record-only capability.</p> <p>Drivers corresponding to other USB audio interfaces on the device, including the usb_as(7D) audio streaming driver or the hid(7D) driver, are plumbed under the USB audio control driver and do not directly interface with user applications.</p> <p>The usb_ac driver supports USB audio class compliant devices with a feature unit. For a list of recommended devices, visit: <a href="http://www.sun.com/io">www.sun.com/io</a>.</p> |
| <b>APPLICATION PROGRAM INTERFACE</b>    | <p>This interface is described in the mixer(7I) and audio(7I) man pages.</p> <p>Applications that open /dev/audio may use the AUDIO_GETDEV ioctl() to determine which audio device is being used. The USB audio driver returns the string "USB Audio" in the name field of the audio_device structure. The version field displays the version number and the config field displays the string "external."</p> <p>The USB audio device provides support for an external speaker and microphone.</p>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Audio Mixer Mode</b>                 | <p>The configuration file /kernel/drv/usb_ac.conf is used to configure the USB audio driver and determines whether the audio mixer is enabled or disabled. See the mixer(7I) manual page for details. You can change the audio mixer mode at any time by using the mixerctl(1) or sdtaudiocontrol(1) applications.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Audio Data Formats</b>               | <p>The USB audio device supports 8-bit <math>\mu</math>-law and A-law, 8-bit linear and 16-bit linear encodings in mono and stereo. With the mixer enabled, a continuous range of sample rates from 8000 to 48000 Hz is supported. With the mixer disabled, the following sample rates are supported: 8000, 9600, 11025, 16000, 18900, 22050, 32000, 33075, 37800, 4410, and 48000 Hz if the device supports these frequencies.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Audio Status Change Notification</b> | <p>As described in the audio(7I) and mixer(7I) man pages, it is possible to request asynchronous notification of changes in the state of an audio device.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>ERRORS</b>                           | <p>If a device is hot-removed while it is active, all subsequent opens will return EIO. All other errors are defined in the audio(7I) man page.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>FILES</b>                            | <p>/kernel/drv/usb_ac<br/>32-bit x86 ELF kernel module</p> <p>/kernel/drv/amd64/usb_ac<br/>64-bit x86 ELF kernel module</p> <p>/kernel/drv/sparcv9/usb_ac<br/>64-bit SPARC ELF kernel module.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## usb\_ac(7D)

/kernel/drv/usb\_ac.conf

USB audio driver configuration file.

/dev/audio

Symlink to the system's primary audio device, not necessarily a USB audio device.

/dev/audioctrl

/dev/audio control device.

/dev/sound/[0-N]

Represents the audio devices on the system and is not necessarily a USB audio device.

/dev/sound/[0-N]ctl

/dev/sound audio control device.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE  | ATTRIBUTE VALUE               |
|-----------------|-------------------------------|
| Architecture    | SPARC, x86, PCI-based systems |
| Availability    | SUNWusb                       |
| Stability level | Evolving                      |

**SEE ALSO** `mixerctl(1)`, `cfgadm_usb(1M)`, `ioctl(2)`, `attributes(5)`, `hid(7D)`, `usba(7D)`, `usb_as(7D)`, `audio(7I)`, `mixer(7I)`, `streamio(7I)`, `usb_ah(7M)`

*Writing Device Drivers*

*Universal Serial Bus Specification 1.0 and 1.1*

*Universal Serial Bus Device Class Definition for Audio Devices, Release 1.0*

*System Administration: Basic Administration*

<http://www.sun.com/io>

**DIAGNOSTICS** In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

Warning: <device path> (usb\_ac<instance num>): Error Message...

Failure to plumb audio streams drivers.

The usb audio streaming driver or the hid driver could not be plumbed under the audio control driver and the device is not usable.

Device was disconnected while open. Data may have been lost.

The device was hot-removed or powered off while it was open and a possible data transfer was in progress. The job was aborted.

Cannot access device. Please reconnect <name>.

There was an error in accessing the device during reconnect. Please reconnect the device.

Device is not identical to the previous one on this port. Please disconnect and reconnect.

A USB audio device was hot-removed while open. A new device was hot-inserted which is not identical to the original USB audio device. Please disconnect the USB device and reconnect the device to the same port.

Busy device has been reconnected.

A device that was hot-removed from a USB port has been re-inserted again.

## NOTES

Upon the initial `open()` of the audio device, the driver resets the data format of the device to the default state of 8-bit, 8Khz, mono u-Law data. If the device is already open and a different audio format is set, this will not be possible on some devices. With the exception of some devices that only support a limited number of sample rates, audio applications should explicitly set the encoding characteristics to match the audio data requirements rather than depend on the default configuration.

The USB audio device will be power managed if the device is idle.

If a USB audio device is hot-removed while active, it prints a console warning message requesting you to put the device back in the same port and informing you that there may be data loss. Hot-removal of an active audio device is strongly discouraged.

Close all applications before hot-removing or hot-inserting a device. If an application is open when a device is hot-removed, inserting the device in a different port will create new `/dev/sound` links but `/dev/audio` will not be affected. Hotplugging an active device is not recommended.

On slower IA machines and with higher frequency sample rates, you may encounter some audio quality problems.

To make a USB audio device the primary audio device (for example: `/dev/audio`), close all audio applications, disconnect all USB audio devices and then simply reconnect the USB audio device. This causes `/dev/audio` to point to the USB audio `/dev/usb/audio` entry.

Most Solaris audio applications and 3rd party audio applications available on Solaris work well with USB audio devices. For details of the application behavior with USB audio devices, visit [www.sun.com/io](http://www.sun.com/io).

## usb\_ah(7M)

| <b>NAME</b>         | usb_ah – USB audio HID STREAMS module                                                                                                                                                                                                                                                                                                                                                                                                              |                |                 |              |                   |              |         |                     |          |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-----------------|--------------|-------------------|--------------|---------|---------------------|----------|
| <b>DESCRIPTION</b>  | The usb_ah STREAMS module enables the USB input control device which is a member of the Human Interface Device (HID) class and provides support for volume change and mute button. The usb_ah module is pushed on top of a HID class driver instance (see hid(7D)) and below an Audio Control class driver instance (see usb_ac(7D)). It translates the HID specific events to the events that are supported by the Solaris audio mixer framework. |                |                 |              |                   |              |         |                     |          |
| <b>FILES</b>        | <p>/kernel/strmod/usb_ah<br/>32-bit ELF kernel STREAMS module. (x86 platform only.)</p> <p>/kernel/strmod/sparcv9/usb_ah<br/>SPARC 64-bit ELF kernel STREAMS module</p>                                                                                                                                                                                                                                                                            |                |                 |              |                   |              |         |                     |          |
| <b>ATTRIBUTES</b>   | See attributes(5) for descriptions of the following attributes:                                                                                                                                                                                                                                                                                                                                                                                    |                |                 |              |                   |              |         |                     |          |
|                     | <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Architecture</td><td>PCI-based systems</td></tr><tr><td>Availability</td><td>SUNWusb</td></tr><tr><td>Interface Stability</td><td>Evolving</td></tr></tbody></table>                                                                                                                                                                       | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Architecture | PCI-based systems | Availability | SUNWusb | Interface Stability | Evolving |
| ATTRIBUTE TYPE      | ATTRIBUTE VALUE                                                                                                                                                                                                                                                                                                                                                                                                                                    |                |                 |              |                   |              |         |                     |          |
| Architecture        | PCI-based systems                                                                                                                                                                                                                                                                                                                                                                                                                                  |                |                 |              |                   |              |         |                     |          |
| Availability        | SUNWusb                                                                                                                                                                                                                                                                                                                                                                                                                                            |                |                 |              |                   |              |         |                     |          |
| Interface Stability | Evolving                                                                                                                                                                                                                                                                                                                                                                                                                                           |                |                 |              |                   |              |         |                     |          |
| <b>SEE ALSO</b>     | <p>mixerctl(1), hid(7D), usba(7D), usb_ac(7D), usb_as(7D), usb_mid(7D), audio(7I), mixer(7I)</p> <p><i>STREAMS Programming Guide</i></p> <p><i>System Administration Guide: Basic Administration</i></p> <p><i>Universal Serial Bus Specification 1.0 and 1.1</i></p> <p><i>Device Class Definition for Human Interface Devices (HID) 1.1</i></p>                                                                                                  |                |                 |              |                   |              |         |                     |          |
| <b>DIAGNOSTICS</b>  | None                                                                                                                                                                                                                                                                                                                                                                                                                                               |                |                 |              |                   |              |         |                     |          |
| <b>NOTES</b>        | If USB audio drivers are not loaded, buttons will not be active.                                                                                                                                                                                                                                                                                                                                                                                   |                |                 |              |                   |              |         |                     |          |

| <b>NAME</b>        | usb_as – USB audio streaming driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                |                 |              |                               |              |         |                 |          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|-----------------|--------------|-------------------------------|--------------|---------|-----------------|----------|
| <b>SYNOPSIS</b>    | sound@unit-address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                |                 |              |                               |              |         |                 |          |
| <b>DESCRIPTION</b> | <p>The usb_as driver is a USBA (Solaris USB Architecture) compliant client driver that supports the <i>USB Audio Class 1.0</i> specification.</p> <p>The usb_as driver processes audio data messages during play and record and sets sample frequency, precision, encoding and other functions on request from the USB audio control driver. See <a href="#">usb_ac(7D)</a>.</p> <p>This driver is plumbed under the USB audio control driver and does not directly interface with the user application.</p> |                |                 |              |                               |              |         |                 |          |
| <b>FILES</b>       | <p>/kernel/drv/usb_as<br/>32-bit x86 ELF kernel module</p> <p>/kernel/drv/amd64/usb_as<br/>64-bit x86 ELF kernel module</p> <p>/kernel/drv/sparcv9/usb_as<br/>64-bit SPARC ELF kernel module</p>                                                                                                                                                                                                                                                                                                             |                |                 |              |                               |              |         |                 |          |
| <b>ATTRIBUTES</b>  | <p>See <a href="#">attributes(5)</a> for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Architecture</td> <td>SPARC, x86, PCI-based systems</td> </tr> <tr> <td>Availability</td> <td>SUNWusb</td> </tr> <tr> <td>Stability level</td> <td>Evolving</td> </tr> </tbody> </table>                                                                                                             | ATTRIBUTE TYPE | ATTRIBUTE VALUE | Architecture | SPARC, x86, PCI-based systems | Availability | SUNWusb | Stability level | Evolving |
| ATTRIBUTE TYPE     | ATTRIBUTE VALUE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                |                 |              |                               |              |         |                 |          |
| Architecture       | SPARC, x86, PCI-based systems                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                |                 |              |                               |              |         |                 |          |
| Availability       | SUNWusb                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                |                 |              |                               |              |         |                 |          |
| Stability level    | Evolving                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                |                 |              |                               |              |         |                 |          |
| <b>SEE ALSO</b>    | <p><a href="#">mixerctl(1)</a>, <a href="#">attributes(5)</a>, <a href="#">usba(7D)</a>, <a href="#">usb_ac(7D)</a>, <a href="#">audio(7I)</a>, <a href="#">mixer(7I)</a>, <a href="#">streamio(7I)</a></p> <p><i>Writing Device Drivers</i></p> <p><i>Universal Serial Bus Specification 1.0 and 1.1</i></p> <p><i>System Administration Guide: Basic Administration</i></p> <p><a href="http://www.sun.com/io">http://www.sun.com/io</a></p>                                                               |                |                 |              |                               |              |         |                 |          |
| <b>DIAGNOSTICS</b> | <p>In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:</p> <p>Warning: &lt;device path&gt; (usb_as&lt;instance num&gt;): Error Message...</p> <p>where &lt;device path&gt; is the physical path to the device in /devices directory.</p> <p>No bandwidth available.<br/>There is no bandwidth available for the isochronous pipe. As a result, no data will be transferred during play and record.</p>                 |                |                 |              |                               |              |         |                 |          |

usb\_as(7D)

Operating a full/high speed audio device on a high speed port is not supported.

The USB software does not currently support full or high speed audio devices connected to an external USB 2.0 hub that is linked to a port of a USB 2.0 host controller. Audio devices must be connected directly to a port of a USB 2.0 controller or to any USB 1.1 port.

Cannot access device. Please reconnect <name>.

There was an error in accessing the device during reconnect. Please reconnect the device.

Device is not identical to the previous one on this port. Please disconnect and reconnect.

A USB audio streaming interface was hot-removed while open. A new device was hot-inserted which is not identical to the original USB audio device. Please disconnect the USB device and reconnect the device to the same port.

**NOTES** The USB audio streaming interface will be power managed if the device is idle.



|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | usbkbm – keyboard STREAMS module for Sun USB Keyboard                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>SYNOPSIS</b>    | <code>open("/dev/kbd", O_RDWR)</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>DESCRIPTION</b> | <p>The usbkbm STREAMS module processes byte streams generated by a keyboard attached to a USB port. USB keyboard is a member of Human Interface Device (HID) Class, and usbkbm only supports the keyboard protocol defined in the specification. Definitions for altering keyboard translation and reading events from the keyboard are in <code>&lt;sys/kbio.h&gt;</code> and <code>&lt;sys/kbd.h&gt;</code>.</p> <p>The usbkbm STREAMS module adheres to the interfaces exported by <a href="#">kb(7M)</a>. Refer to the DESCRIPTION section of <a href="#">kb(7M)</a> for a discussion of the keyboard translation modes and the IOCTL section for the supported <code>ioctl(2)</code> requests.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>IOCTLS</b>      | <p>USB Keyboard usbkbm returns different values for the following <code>ioctl(2)</code> requests than <a href="#">kb(7M)</a>:</p> <p><b>KIOCTYPE</b></p> <p>This <code>ioctl()</code> returns a new keyboard type defined for the USB keyboard. All types are listed below:</p> <pre>KB_SUN3   Sun Type 3 keyboard KB_SUN4   Sun Type 4 keyboard KB_ASCII  ASCII terminal masquerading as keyboard KB_PC     Type 101 PC keyboard KB_USB    USB keyboard</pre> <p>The USB keyboard type is <code>KB_USB</code>; usbkbm will return <code>KB_USB</code> in response to the <code>KIOCTYPE</code> <code>ioctl</code>.</p> <p><b>KIOCLAYOUT</b></p> <p>The argument is a pointer to an <code>int</code>. The layout code specified by the <code>bCountryCode</code> value returned in the HID descriptor is returned in the <code>int</code> pointed to by the argument. The country codes are defined in 6.2.1 of the HID 1.0 specifications.</p> <p><b>KIOCCMD</b></p> <p><code>KBD_CMD_CLICK/KBD_CMD_NOCLICK</code></p> <p>The <a href="#">kb(7M)</a> indicates that inappropriate commands for particular keyboards are ignored. Because clicking is not supported on the USB keyboard, usbkbm ignores this command</p> <p><code>KBD_CMD_SETLED</code></p> <p>Set keyboard LEDs. Same as <a href="#">kb(7M)</a>.</p> <p><code>KBD_CMD_GETLAYOUT</code></p> <p>The country codes defined in 6.2.1 of the HID 1.0 specification are returned.</p> <p><code>KBD_CMD_BELL/KBD_CMD_NOBELL</code></p> <p>This command is supported although the USB keyboard does not have a buzzer. The request for the bell is rerouted.</p> <p><code>KBD_CMD_RESET</code></p> <p>There is no notion of resetting the keyboard as there is for the type4 keyboard. usbkbm ignores this command and does not return an error.</p> |

usbkbm(7M)

**FILES** /kernel/strmod/usbkbm  
32-bit ELF kernel STREAMS module (x86 platform only)  
  
/kernel/strmod/sparcv9/usbkbm  
SPARC 64-bit ELF kernel STREAMS module

**ATTRIBUTES** See attributes(5) for a description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE   |
|----------------|-------------------|
| Architecture   | PCI-based systems |
| Availability   | SUNWusb           |

**SEE ALSO** dumpkeys(1), kbd(1), loadkeys(1), ioctl(2), keytables(4), attributes(5), hid(7D), [usba\(7D\)](#), termio(7I), [kb\(7M\)](#)

*STREAMS Programming Guide*

*System Administration Guide: Basic Administration*

<http://www.sun.com/io>

**DIAGNOSTICS** None

|                    |                                                                                                                                                                                                                                                                                                                        |                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| <b>NAME</b>        | usb_mid – USB Multi Interface Driver                                                                                                                                                                                                                                                                                   |                                |
| <b>SYNOPSIS</b>    | device@unit-address                                                                                                                                                                                                                                                                                                    |                                |
| <b>DESCRIPTION</b> | The usb_mid driver is a USBA (Solaris Universal Serial Bus Architecture) compliant nexus driver that binds to device level nodes of a composite (multi interface) device if no vendor or class specific driver is available. The usb_mid driver attempts to bind drivers to each of the composite device's interfaces. |                                |
| <b>FILES</b>       | /kernel/drv/usb_mid                                                                                                                                                                                                                                                                                                    | 32-bit x86 ELF kernel module   |
|                    | /kernel/drv/amd64/usb_mid                                                                                                                                                                                                                                                                                              | 64-bit x86 ELF kernel module   |
|                    | /kernel/drv/sparcv9/usb_mid                                                                                                                                                                                                                                                                                            | 64-bit SPARC ELF kernel module |
| <b>ATTRIBUTES</b>  | See attributes(5) for descriptions of the following attributes:                                                                                                                                                                                                                                                        |                                |

| ATTRIBUTE TYPE | ATTRIBUTE VALUE               |
|----------------|-------------------------------|
| Architecture   | SPARC, x86, PCI-based systems |
| Availability   | SUNWusb                       |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SEE ALSO</b>    | cfgadm_usb(1M), attributes(5), usba(7D)<br><i>Writing Device Drivers</i><br><i>Universal Serial Bus Specification 1.0 and 1.1</i><br><i>System Administration Guide: Basic Administration</i><br><a href="http://www.sun.com/io">http://www.sun.com/io</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>DIAGNOSTICS</b> | In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:<br><br>Warning: <device path> (usb_mid<instance number>): Error Message...<br><br>Cannot access <device>. Please reconnect.<br>This device has been disconnected because a device other than the original one has been inserted. The driver informs you of this fact by displaying the name of the original device.<br><br>Device not identical to the previous one on this port. Please disconnect and reconnect.<br>Same condition as described above; however in this case, the driver is unable to identify the original device with a name string.<br><br>Please disconnect and reconnect this device.<br>A hotplug of the device is needed before it can be restored. |

## usbms(7M)

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | usbms – USB mouse STREAMS module                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>SYNOPSIS</b>    | <pre>#include &lt;sys/vuid_event.h&gt; #include &lt;sys/vuid_wheel.h&gt; #include &lt;sys/msio.h&gt; #include &lt;sys/msreg.h&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>DESCRIPTION</b> | <p>The usbms STREAMS module processes byte streams generated by a USB mouse. A USB mouse is a member of the Human Interface Device (HID) class and the usbms module supports only the mouse boot protocol defined in the HID specification. The usbms module must be pushed on top of the HID class driver (see hid(7D)). In the VUID_FIRM_EVENT mode, the usbms module translates packets from the USB mouse into Firm events. The Firm event structure is defined in &lt;sys/vuid_event.h&gt;. The STREAMS module state is initially set to raw or VUID_NATIVE mode which performs no message processing. See the HID 1.0 specification for the raw format of the mouse packets. To initiate mouse protocol conversion to Firm events, change the state to VUID_FIRM_EVENT.</p>                                                                                                                                                                                                                                                                                                                                                   |
| <b>IOCTLS</b>      | <p><b>VUIDGFORMAT</b> This option returns the current state of the STREAMS module. The state of the usbms STREAMS module may be either VUID_NATIVE (no message processing) or VUID_FIRM_EVENT (convert to Firm events).</p> <p><b>VUIDSFORMAT</b> The argument is a pointer to an int. Set the state of the STREAMS module to the int pointed to by the argument.</p> <pre>typedef struct vuid_addr_probe {     short base; /* default vuid device addr directed too */     union {         short next; /* next addr for default when VUIDSADDR */         short current; /* current addr of default when VUIDGADDR */     } data; } Vuid_addr_probe;</pre> <p><b>VUIDSADDR</b> The argument is a pointer to a Vuid_addr_probe structure. VUIDSADDR sets the virtual input device segment address indicated by base to next.</p> <p>If base does not equal VKEY_FIRST, ENODEV is returned.</p> <p><b>VUIDGADDR</b> The argument is a pointer to a Vuid_addr_probe structure. Return the address of the virtual input device segment indicated by base to current.</p> <p>If base does not equal VKEY_FIRST, ENODEV is returned.</p> |

**VOIDGWHEELCOUNT**

This ioctl takes a pointer to an integer as argument and sets the value of the integer to the number of wheels available on this device. This ioctl returns 1 if wheel(s) are present and zero if no wheels are present.

**VOIDGWHEELINFO**

This command returns static information about the wheel that does not change while a device is in use. Currently the only information defined is the wheel orientation which is either `VOID_WHEEL_FORMAT_VERTICAL` or `VOID_WHEEL_FORMAT_HORIZONTAL`. If the module cannot distinguish the orientation of the wheel or the wheel is of some other format, the format is set to `VOID_WHEEL_FORMAT_UNKNOWN`.

```
typedef struct {
 int vers;
 int id;
 int format;
} wheel_info;
```

The ioctl takes a pointer to "wheel\_info" structure with the "vers" set to the current version of the "wheel\_info" structure and "id" set to the id of the wheel for which the information is desired.

**VOIDSWHEELSTATE/VOIDGWHEELSTATE**

`VOIDSWHEELSTATE` sets the state of the wheel to that specified in the stateflags. `VOIDGWHEELSTATE` returns the current state settings in the stateflags field.

stateflags is an OR'ed set of flag bits. The only flag currently defined is `VOID_WHEEL_STATE_ENABLED`.

When stateflags is set to `VOID_WHEEL_STATE_ENABLED` the module converts motion of the specified wheel into VOID events and sends those up stream.

Wheel events are enabled by default.

Applications that want to change the stateflags should first get the current stateflags and then change only the bit they want.

```
typedef struct {
 int vers;
 int id;
 uint32_t stateflags;
} wheel_state;
```

These ioctls take a pointer to "wheel\_state" as an argument with the "vers" and "id" members filled in. These members have the same meaning as that for 'VOIDGWHEEL INFO' ioctl.

`ioctl()` requests for changing and retrieving mouse parameters use the `Ms_parms` structure:

```
typedef struct {
 int jitter_thresh;
 int speed_low;
```

## usbms(7M)

```
 int speed_limit;
 } Ms_params;
```

`jitter_thresh` is the "jitter threshold" of the mouse. Motions fewer than `jitter_thresh` units along both axes are accumulated and then sent up the stream after 1/12 second.

`speed_low` indicates whether extremely large motions are to be ignored. If it is 1, a "speed limit" is applied to mouse motions. Motions along either axis of more than `speed_limit` units are discarded.

**MSIOGETPARMS** The argument is a pointer to a `Ms_params` structure. The `usbms` module parameters are returned in the structure.

**MSIOSETPARMS** The argument is a pointer to a `Ms_params` structure. The `usbms` module parameters are set according to the values in the structure.

**FILES** /kernel/strmod/usbms  
32-bit ELF kernel STREAMS module (x86 platform only.)  
  
/kernel/strmod/sparcv9/usbms  
SPARC 64-bit ELF kernel STREAMS module

**ATTRIBUTES** See `attributes(5)` for a description of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE   |
|----------------|-------------------|
| Architecture   | PCI-based systems |
| Availability   | SUNWusb           |

**SEE ALSO** `ioctl(2)`, `attributes(5)`, `hid(7D)`, `usba(7D)`

*System Administration Guide: Basic Administration*

<http://www.sun.com/io>

**DIAGNOSTICS** None

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>              | usbprn – USB printer class driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>SYNOPSIS</b>          | <pre>#include &lt;sys/usb/clients/printer/usb_printer.h&gt; #include &lt;sys/ecppio.h&gt; usbprn@unit-address</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>DESCRIPTION</b>       | <p>The usbprn driver is a USBA (Solaris USB Architecture) compliant client driver that supports the <i>USB Printer Class 1.0</i> specification. The usbprn driver supports a subset of the <a href="#">ecpp(7D)</a> parallel port driver functionality. However, unlike the STREAMS-based ecpp driver, usbprn is a character driver.</p> <p>The usbprn driver supports all USB printer-class compliant printers. For a list of recommended printers and USB parallel printer adapters, visit <a href="http://www.sun.com/io">http://www.sun.com/io</a>.</p> <p>The usbprn driver includes support for communicating with many different printers. To use these printers, it may be necessary to install and configure additional format conversion packages available in the Solaris distribution. Configuration of these conversion packages under the Solaris printing system can be simplified through the use of the <a href="#">printmgr(1M)</a>. This tool allows selection of printer manufacturer/model information while creating a print queue. For USB connected printers, it attempts to pre-select the the manufacturer and model information based on the 1284 device id supplied by the printer.</p> |
| <b>DEFAULT OPERATION</b> | <p>With certain minor exceptions (outlined in the Notes sections below), the usbprn driver supports a subset of the <a href="#">ecpp(7D)</a> ioctl interfaces:</p> <p>Configuration variables are set to their default values each time the USB printer device is attached. The <code>write_timeout</code> period (defined in the <code>ECPPIOC_SETPARMS</code> ioctl description below) is set to 90 seconds. The mode is set to centronics mode (<code>ECPPIOC_SETPARMS</code> ioctl and read through the <code>ECPPIOC_GETPARMS</code> ioctl. Each time the USB printer device is opened, the device is marked as busy and all further opens will return <code>EBUSY</code>. Once the device is open, applications can write to the device and the driver can send data and obtain device id and status.</p> <p><b>Note</b> – Unlike the <a href="#">ecpp(7D)</a> driver, usbprn resets configuration variables to their default values with each <code>attach(9E)</code>. (The <a href="#">ecpp(7D)</a> driver resets configuration variables with each <code>open(2)</code>.)</p>                                                                                                                              |
| <b>WRITE OPERATION</b>   | <p>A <code>write(2)</code> operation returns the number of bytes successfully written to the device. If a failure occurs while a driver is transferring data to printer, the contents of the status bits are captured at the time of the error and can be retrieved by the application program using the <code>ECPPIOC_GETERR</code> <code>ioctl(2)</code> call. The captured status information is overwritten each time an <code>ECPPIOC_TESTIO</code> <code>ioctl(2)</code> occurs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>IOCTLS</b>            | <p>The usbprn driver supports <code>prnio(7I)</code> interfaces. Note that the <code>PRNIOC_RESET</code> command has no effect on USB printers.</p> <p>The following <code>ioctl(2)</code> calls are supported for backward compatibility and are not recommended for new applications.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## usbprn(7D)

### ECPPIOC\_GETPARMS

Gets current transfer parameters. The argument is a pointer to `struct ecpp_transfer_parms`. If parameters are not configured after the device is opened, the structure will be set to its default configuration.

**Note** – Unlike the `ecpp(7D)` driver, only the `ECPP_CENTRONICS` mode is currently supported in `usbprn`.

### ECPPIOC\_SETPARMS

Sets transfer parameters. The argument is a pointer to a `struct ecpp_transfer_parms`. If a parameter is out of range, `EINVAL` is returned. If the peripheral or host device cannot support the requested mode, `EPROTONOSUPPORT` is returned.

The transfer parameters structure is defined in `<sys/ecppio.h>`:

```
struct ecpp_transfer_parms {
 int write_timeout;
 int mode;
};
```

The `write_timeout` field, which specifies how long the driver will take to transfer 8192 bytes of data to the device, is set to a default value of 90 seconds. The `write_timeout` field must be greater than one second and less than 300 seconds (five minutes.)

**Note** – Unlike the `ecpp(7D)` driver, only the `ECPP_CENTRONICS` mode is currently supported in `usbprn`. Also, the semantics of `write_timeout` in `usbprn` differ from `ecpp(7D)`. Refer to `ecpp(7D)` for information.

### BPPIOC\_TESTIO

Tests the transfer readiness of a print device and checks status bits to determine if a `write(2)` will succeed. If status bits are set, a transfer will fail. If a transfer will succeed, zero is returned. If a transfer fails, the driver returns `EIO` and the state of the status bits are captured. The captured status can be retrieved using the `BPPIOC_GETERR ioctl(2)` call. `BPPIOC_TESTIO` and `BPPIOC_GETERR` are compatible to the `ioctls` specified in `bpp(7D)`.

**Note** – Unlike the `ecpp(7D)` driver, only the `ECPP_CENTRONICS` mode is currently supported in `usbprn`. Additionally, `bus_error` and `timeout_occurred` fields are not used in the `usbprn` interface. (In `ecpp(7D)`, `timeout_occurred` is used.)

### BPPIOC\_GETERR

Get last error status. The argument is a pointer to a `struct bpp_error_status`. This structure indicates the status of all the appropriate status bits at the time of the most recent error condition during a `write(2)` call, or the status of the bits at the most recent `BPPIOC_TESTIO ioctl(2)` call.

```
struct bpp_error_status {
 char timeout_occurred; /* not used */
 char bus_error; /* not used */
 uchar_t pin_status; /* status of pins which
 /* could cause error */
};
```



The `pin_status` field indicates possible error conditions. The error status structure `bpp_error_status` is defined in the include file `<sys/bpp_io.h>`. The valid bits for `pin_status` can be `BPP_ERR_ERR`, `BPP_SLCT_ERR`, and `BPP_PE_ERR`. A set bit indicates that the associated pin is asserted.

**Note** – Unlike the `ecpp(7D)` driver, only the `ECPP_CENTRONICS` mode is currently supported in `usbprn`. Additionally, the `bus_error` and `timeout_occurred` fields are not used in the `usbprn` interface. (In `ecpp(7D)`, `timeout_occurred` is used.) Unlike `ecpp(7D)`, the `BPP_BUSY_ERR` status bit is not supported by USB printers.

#### ECPPIOC\_GETDEVID

Gets the IEEE 1284 device ID from the peripheral. The argument is a pointer to a structure `ecpp_device_id`. Applications should set `mode` to `ECPP_CENTRONICS`. If another mode is used, the driver will return `EPROTONOSUPPORT`. `len` is the length of the buffer pointed to by `addr`. `rlen` is the actual length of the device ID string returned from the peripheral. If the returned `rlen` is greater than `len`, the application should call `ECPPIOC_GETDEVID` a second time with a buffer length equal to `rlen`.

The 1284 device ID structure:

```
struct ecpp_device_id {
 int mode; /* mode to use for reading device id */
 int len; /* length of buffer */
 int rlen; /* actual length of device id string */
 char *addr; /* buffer address */
}
```

**Note** – Unlike `ecpp(7D)`, only the `ECPP_CENTRONICS` mode is currently supported in `usbprn`.

## READ OPERATION ERRORS

The read operation is not supported and returns `EIO`.

|                     |                                                                                                                                                                                                                                                                                                |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>EBUSY</code>  | The device has been opened and another open is attempted. An attempt has been made to unload the driver while one of the units is open.                                                                                                                                                        |
| <code>EINVAL</code> | An unsupported IOCTL has been received. A <code>ECPPIOC_SETPARMS ioctl(2)</code> is attempted with an out of range value in the <code>ecpp_transfer_parms</code> structure.                                                                                                                    |
| <code>EIO</code>    | The driver has received an unrecoverable device error, or the device is not responding, or the device has stalled when attempting an access. A <code>write(2)</code> or <code>ioctl(2)</code> did not complete due to a peripheral access. A <code>read(2)</code> system call has been issued. |
| <code>ENXIO</code>  | The driver has received an <code>open(2)</code> request for a unit for which the attach failed.                                                                                                                                                                                                |
| <code>ENODEV</code> | The driver has received an <code>open(2)</code> request for a device that has been disconnected.                                                                                                                                                                                               |

## usbprn(7D)

EPROTONOSUPPORT      The driver has received a ECPIOC\_SETPARMS ioctl(2) for a mode argument other than ECPP\_CENTRONICS in the ecpp\_transfer\_parms structure.

**FILES**

- /kernel/drv/usbprn  
32-bit x86 ELF kernel module
- /kernel/drv/amd64/usbprn  
64-bit x86 ELF kernel module
- /kernel/drv/sparcv9/usbprn  
64-bit SPARC ELF kernel module
- /dev/printers/*n*  
Character special files

**ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE               |
|----------------|-------------------------------|
| Architecture   | SPARC, x86, PCI-based systems |
| Availability   | SUNWusb                       |

**SEE ALSO** cfgadm\_usb(1M), printmgr(1M), ioctl(2), open(2), read(2), write(2), attributes(5), [bpb\(7D\)](#), [ecpp\(7D\)](#), [usba\(7D\)](#)prnio(7I), attach(9E)

*Writing Device Drivers*

*Universal Serial Bus Specification 1.0 and 1.1*

*USB Device Class Definition for Printing Devices 1.0*

*System Administration Guide: Basic Administration*

<http://www.sun.com/io>

**DIAGNOSTICS** In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

Warning: <device path> (usbprn<instance num>): Error Message...

Device was disconnected while open. Data may have been lost.

The device has been hot-removed or powered off while it was open and a possible data transfer was in progress. The job may be aborted.

Cannot access <device>. Please reconnect.

There was an error in accessing the printer during reconnect. Please reconnect the device.

Device is not identical to the previous one on this port. Please disconnect and reconnect.

A USB printer was hot-removed while open. A new device was hot-inserted which is not identical to the original USB printer. Please disconnect the USB device and reconnect the printer to the same port.

Printer has been reconnected but data may have been lost.

The printer that was hot-removed from its USB port has been re-inserted again to the same port. It is available for access but the job that was running prior to the hot-removal may be lost.

**NOTES** The USB printer will be power managed if the device is closed.

If a printer is hot-removed before a job completes, the job is terminated and the driver will return EIO. All subsequent opens will return `ENODEV`. If a printer is hot-removed, an LP reconfiguration may not be needed if a printer is re-inserted on the same port. If re-inserted on a different port, an LP reconfiguration may be required.

The USB Parallel Printer Adapter is not hotpluggable. The printer should be connected to USB Parallel Printer Adapter before plugging the USB cable into host or hub port and should be removed only after disconnecting the USB cable of USB Parallel Printer Adapter from the host or hub port.

## usbser\_edge(7D)

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                              | usbser_edge – Digi Edgeport USB to serial converter driver                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>SYNOPSIS</b>                          | <pre>#include &lt;fcntl.h&gt; #include &lt;sys/termios.h&gt; usbser_edge@unit</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>DESCRIPTION</b>                       | <p>The <code>usbser_edge</code> driver is a loadable STREAMS and USBA (Solaris USB architecture) compliant client driver that provides basic asynchronous communication support for Digi Edgeport USB-to-serial converters. Supported devices include Edgeport/1, Edgeport/2, Edgeport/21, Edgeport/4, Edgeport/421, Edgeport/8, and Edgeport/416. Serial device streams are built with appropriate modules that are pushed atop the <code>usbser_edge</code> driver by the <code>autopush(1M)</code> facility.</p> <p>The <code>usbser_edge</code> module supports the <code>termio(7I)</code> device control functions specified by flags in the <code>c_cflag</code> word of the <code>termios</code> structure, and by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code> and <code>INPCK</code> flags in the <code>c_iflag</code> word of the <code>termios</code> structure. All other <code>termio(7I)</code> functions must be performed by STREAMS modules pushed atop the driver. When a device is opened, the <code>ldterm(7M)</code> and <code>ttcompat(7M)</code> STREAMS modules are automatically pushed on top of the stream, providing the standard <code>termio(7I)</code> interface.</p> <p>Use device logical names <code>/dev/term/[0-9]*</code> to access the serial ports. These names are typically used to provide a logical access point for a dial-in line that is used with a modem.</p> <p>To allow a single tty line to be connected to a modem and used for incoming and outgoing calls, a special feature is available that is controlled by the minor device number. By accessing through device logical name <code>/dev/cua/[0-9]*</code>, you can open a port without the Carrier Detect signal being asserted, either through hardware or an equivalent software mechanism. These devices are commonly known as dial-out lines.</p> <p>Unlike onboard serial ports, the <code>usbser_edge</code> ports cannot serve as a local serial console.</p> |
| <b>APPLICATION PROGRAMMING INTERFACE</b> | <p>A dial-in line can be opened only if the corresponding dial-out line is closed. A blocking <code>/dev/term</code> <code>open</code> waits until the <code>/dev/cua</code> line is closed (which drops Data Terminal Ready, after which Carrier Detect usually drops as well) and carrier is detected again. A non-blocking <code>/dev/term</code> <code>open</code> returns an error if the <code>/dev/cua</code> is open.</p> <p>If the <code>/dev/term</code> line is opened successfully (usually only when carrier is recognized on the modem), the corresponding <code>/dev/cua</code> line cannot be opened. This allows a modem and port to be used for dial-in (by enabling the line for login in <code>/etc/inittab</code>) or dial-out (by <code>tip(1)</code>, or <code>uucp(1C)</code>) when no one is logged in on the line.</p> <p>Device hot-removal is functionally equivalent to modem disconnect event, as defined in <code>termio(7I)</code>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>IOCTLS</b>                            | The <code>usbser_edge</code> driver supports the standard set of <code>termio(7I)</code> <code>ioctl</code> calls.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Input and output line speeds can be set to the following baud rates: 0, 50, 75, 110, 134, 150, 200, 300, 600, 1200, 1800, 2400, 4800, 9600, 19200, 38400, 57600, 115200, or 230400. Input and output line speeds cannot be set independently; for example, when the output speed is set, the input speed is automatically set to the same speed.

**ERRORS** An `open()` fails under the following conditions:

|       |                                                                                                                                                                                                                      |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENXIO | The unit being opened does not exist.                                                                                                                                                                                |
| EBUSY | The <code>/dev/cua</code> (dial-out) device is being opened while the <code>/dev/term</code> (dial-in device) is open, or the dial-in device is being opened with a no-delay open while the dial-out device is open. |
| EBUSY | The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call.                                                                                                              |
| EIO   | USB device I/O error.                                                                                                                                                                                                |

**FILES**

- `/kernel/drv/usbser_edge`  
32-bit x86 ELF kernel module
- `/kernel/drv/amd64/usbser_edge`  
64-bit x86 ELF kernel module
- `/kernel/drv/sparcv9/usbser_edge`  
64-bit SPARC ELF kernel module
- `/dev/cua/[0-9]*`  
dial-out tty lines
- `/dev/term/[0-9]*`  
dial-in tty lines

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE               |
|----------------|-------------------------------|
| Architecture   | SPARC, x86, PCI-based systems |
| Availability   | SUNWuedg                      |

**SEE ALSO** `strconf(1)`, `tip(1)`, `uucp(1C)`, `autopush(1M)`, `ioctl(2)`, `open(2)`, `termios(3C)`, `attributes(5)`, `usba(7D)`, `termio(7I)`, `ldterm(7M)`, `ttcompat(7M)`

**DIAGNOSTICS** In addition to being logged, the following messages may appear on the system console. All messages are formatted in the following manner:

```
Warning: <device path> (usbser_edge<instance num>): Error Message...
```

Device was disconnected while open. Data may have been lost.

The device was hot-removed or powered off while it was open and a possible data transfer was in progress.

usbser\_edge(7D)

Device is not identical to the previous one on this port. Please disconnect and reconnect.

The device was hot-removed while open. A new device was hot-inserted which is not identical to the original device. Please disconnect the device and reconnect the original device to the same port.

Device has been reconnected, but data may have been lost.

The device that was hot-removed from its USB port has been re-inserted again to the same port. It is available for access but data from a previous transfer may be lost.

Cannot access <device>. Please reconnect.

This device has been disconnected because a device other than the original one has been inserted. The driver informs you of this fact by displaying the name of the original device.

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|----------------------------------------------------------------------------------------------------------------|---------------------------|------------------------------------------------------------------------|----------------------------|---------------------------------------------------------|------------------------|--------------------------------------------------------------------------------------------------------|
| <b>NAME</b>                | uscsi – user SCSI command interface                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |
| <b>SYNOPSIS</b>            | <pre>#include &lt;sys/scsi/impl/uscsi.h&gt;  ioctl(int <i>fildev</i>, int <i>request</i>, struct uscsi_cmd *<i>cmd</i>);</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |
| <b>DESCRIPTION</b>         | <p>The <code>uscsi</code> command is very powerful and somewhat dangerous; therefore it has some permission restrictions. See <b>WARNINGS</b> for more details.</p> <p>Drivers supporting this <code>ioctl(2)</code> provide a general interface allowing user-level applications to cause individual SCSI commands to be directed to a particular SCSI or ATAPI device under control of that driver. The <code>uscsi</code> command is supported by the <code>sd</code> driver for SCSI disks and ATAPI CD-ROM drives, and by the <code>st</code> driver for SCSI tape drives. <code>uscsi</code> may also be supported by other device drivers; see the specific device driver manual page for complete information.</p> <p>Applications must not assume that all Solaris disk device drivers support the <code>uscsi</code> <code>ioctl</code> command. The SCSI command may include a data transfer to or from that device, if appropriate for that command. Upon completion of the command, the user application can determine how many bytes were transferred and the status returned by the device. Also, optionally, if the command returns a Check Condition status, the driver will automatically issue a Request Sense command and return the sense data along with the original status. See the <code>USCSI_RQENABLE</code> flag below for this Request Sense processing. The <code>uscsi_cmd</code> structure is defined in <code>&lt;sys/scsi/impl/uscsi.h&gt;</code> and includes the following members:</p> <pre>int uscsi_flags;                                /* read, write, etc. see below */ short uscsi_status;                            /* resulting status */ short uscsi_timeout;                          /* Command Timeout */ caddr_t uscsi_cdb                             /* CDB to send to target */ caddr_t uscsi_bufaddr;                        /* i/o source/destination */ size_t uscsi_buflen;                          /* size of i/o to take place*/ size_t uscsi_resid;                           /* resid from i/o operation */ uchar_t uscsi_cdblen;                         /* # of valid CDB bytes */ uchar_t uscsi_rqlen;                          /* size of uscsi_rqbuf */ uchar_t uscsi_rqstatus;                       /* status of request sense cmd */ uchar_t uscsi_rqresid;                        /* resid of request sense cmd */ caddr_t uscsi_rqbuf;                          /* request sense buffer */ void *uscsi_reserved_5;                       /* Reserved for future use */</pre> <p>The fields of the <code>uscsi_cmd</code> structure have the following meanings:</p> <table border="0"> <tr> <td style="vertical-align: top;"><code>uscsi_flags</code></td> <td>The I/O direction and other details of how to carry out the SCSI command. Possible values are described below.</td> </tr> <tr> <td style="vertical-align: top;"><code>uscsi_status</code></td> <td>The SCSI status byte returned by the device is returned in this field.</td> </tr> <tr> <td style="vertical-align: top;"><code>uscsi_timeout</code></td> <td>Time in seconds to allow for completion of the command.</td> </tr> <tr> <td style="vertical-align: top;"><code>uscsi_cdb</code></td> <td>A pointer to the SCSI CDB (command descriptor block) to be transferred to the device in command phase.</td> </tr> </table> | <code>uscsi_flags</code> | The I/O direction and other details of how to carry out the SCSI command. Possible values are described below. | <code>uscsi_status</code> | The SCSI status byte returned by the device is returned in this field. | <code>uscsi_timeout</code> | Time in seconds to allow for completion of the command. | <code>uscsi_cdb</code> | A pointer to the SCSI CDB (command descriptor block) to be transferred to the device in command phase. |
| <code>uscsi_flags</code>   | The I/O direction and other details of how to carry out the SCSI command. Possible values are described below.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |
| <code>uscsi_status</code>  | The SCSI status byte returned by the device is returned in this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |
| <code>uscsi_timeout</code> | Time in seconds to allow for completion of the command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |
| <code>uscsi_cdb</code>     | A pointer to the SCSI CDB (command descriptor block) to be transferred to the device in command phase.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                          |                                                                                                                |                           |                                                                        |                            |                                                         |                        |                                                                                                        |

## uscsi(7I)

|                               |                                                                                                                                                                                              |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>uscsi_bufaddr</code>    | The user buffer containing the data to be read from or written to the device.                                                                                                                |
| <code>uscsi_buflen</code>     | The length of <code>uscsi_bufaddr</code> .                                                                                                                                                   |
| <code>uscsi_resid</code>      | If a data transfer terminates without transferring the entire requested amount, the remainder, or residue, is returned in this field.                                                        |
| <code>uscsi_cdblen</code>     | The length of the SCSI CDB to be transferred to the device in command phase.                                                                                                                 |
| <code>uscsi_rqlen</code>      | The length of <code>uscsi_rqbuf</code> , the application's Request Sense buffer.                                                                                                             |
| <code>uscsi_rqstatus</code>   | The SCSI status byte returned for the Request Sense command executed automatically by the driver in response to a Check Condition status return.                                             |
| <code>uscsi_rqresid</code>    | The residue, or untransferred data length, of the Request Sense data transfer (the number of bytes, less than or equal to <code>uscsi_rqlen</code> , which were not filled with sense data). |
| <code>uscsi_rqbuf</code>      | Points to a buffer in application address space to which the results of an automatic Request Sense command are written.                                                                      |
| <code>uscsi_reserved_5</code> | Reserved for future use.                                                                                                                                                                     |

The `uscsi_flags` field defines the following:

```
USCSI_WRITE /* send data to device */
USCSI_SILENT /* no error messages */
USCSI_DIAGNOSE /* fail if any error occurs */
USCSI_ISOLATE /* isolate from normal commands */
USCSI_READ /* get data from device */
USCSI_ASYNC /* set bus to asynchronous mode */
USCSI_SYNC /* return bus to sync mode if possible */
USCSI_RESET /* reset target */
USCSI_RESET_ALL /* reset all targets */
USCSI_RQENABLE /* enable request sense extensions */
USCSI_RENEGOT /* renegotiate wide/sync on next I/O */
```

The `uscsi_flags` bits have the following interpretation:

|                           |                                                                                                                          |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <code>USCSI_WRITE</code>  | Data will be written from the initiator to the target.                                                                   |
| <code>USCSI_SILENT</code> | The driver should not print any console error messages or warnings regarding failures associated with this SCSI command. |



|                 |                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USCSI_DIAGNOSE  | The driver should not attempt any retries or other recovery mechanisms if this SCSI command terminates abnormally in any way.                                                                                                                                                                                                                                                                                            |
| USCSI_ISOLATE   | This SCSI command should not be executed with other commands.                                                                                                                                                                                                                                                                                                                                                            |
| USCSI_READ      | Data will be read from the target to the initiator.                                                                                                                                                                                                                                                                                                                                                                      |
| USCSI_ASYNC     | Set the SCSI bus to asynchronous mode before running this command.                                                                                                                                                                                                                                                                                                                                                       |
| USCSI_SYNC      | Set the SCSI bus to synchronous mode before running this command.                                                                                                                                                                                                                                                                                                                                                        |
| USCSI_RESET     | Send a SCSI Bus Device Reset Message to this target.                                                                                                                                                                                                                                                                                                                                                                     |
| USCSI_RESET_ALL | Cause a SCSI Bus Reset on the bus associated with this target.                                                                                                                                                                                                                                                                                                                                                           |
| USCSI_RQENABLE  | Enable Request Sense extensions. If the user application is prepared to receive sense data, this bit must be set, the fields <code>uscsi_rqbuf</code> and <code>uscsi_rqbuflen</code> must be non-zero, and the <code>uscsi_rqbuf</code> must point to memory writable by the application.                                                                                                                               |
| USCSI_RENEGOT   | Tells USCSI to renegotiate wide mode and synchronous transfer speed before the transmitted SCSI command is executed. This flag in effects tells the target driver to pass the <code>FLAG_RENEGOTIATE_WIDE_SYNC</code> flag in the SCSI packet before passing the command to an adapter driver for transport.<br><br>See the <code>scsi_pkt(9S)</code> flag <code>FLAG_RENEGOTIATE_WIDE_SYNC</code> for more information. |

**IOCTLS** The `ioctl` supported by drivers providing the `uscsi` interface is:

|          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USCSICMD | The argument is a pointer to a <code>uscsi_cmd</code> structure. The SCSI device addressed by that driver is selected, and given the SCSI command addressed by <code>uscsi_cdb</code> . If this command requires a data phase, the <code>uscsi_buflen</code> and <code>uscsi_bufaddr</code> fields must be set appropriately; if data phase occurs, the <code>uscsi_resid</code> is returned as the number of bytes not transferred. The status of the command, as returned by the device, is returned in the <code>uscsi_status</code> field. If the command terminates with Check Condition status, and Request Sense is enabled, the sense data itself is returned in <code>uscsi_rqbuf</code> . The <code>uscsi_rqresid</code> provides the residue of the Request Sense data transfer. |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

**ERRORS** EINVAL A parameter has an incorrect, or unsupported, value.

uscsi(7I)

- EIO An error occurred during the execution of the command.
- EPERM A process without root credentials tried to execute the USCSICMD ioctl.
- EFAULT The `uscsi_cmd` itself, the `uscsi_cdb`, the `uscsi_buf`, or the `uscsi_rqbuf` point to an invalid address.

**ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

| ATTRIBUTE TYPE | ATTRIBUTE VALUE |
|----------------|-----------------|
| Availability   | SUNWhea         |

**SEE ALSO** `ioctl(2)`, `attributes(5)`, `sd(7D)`, `st(7D)`

*ANSI Small Computer System Interface-2 (SCSI-2)*

**WARNINGS** The `uscsi` command is very powerful, but somewhat dangerous, and so its use is restricted to processes running as root, regardless of the file permissions on the device node. The device driver code expects to own the device state, and `uscsi` commands can change the state of the device and confuse the device driver. It is best to use `uscsi` commands only with no side effects, and avoid commands such as Mode Select, as they may cause damage to data stored on the drive or system panics. Also, as the commands are not checked in any way by the device driver, any block may be overwritten, and the block numbers are absolute block numbers on the drive regardless of which slice number is used to send the command.

The `uscsi` interface is not recommended for very large data transfers (typically more than 16MB). If the requested transfer size exceeds the maximum transfer size of the DMA engine, it will not be broken up into multiple transfers and DMA errors may result.

usoc(7D)

- NAME** usoc – universal serial optical controller for Fibre Channel arbitrated loop (SOC+) device driver
- DESCRIPTION** The Fibre Channel adapter is an SBus card that implements two full duplex Fibre Channel interfaces. Each interface can connect to a Fibre Channel arbitrated loop (FC-AL). The usoc device driver is a nexus driver and implements portions of the FC-2 and FC-4 layers of FC-AL.
- FILES** /kernel/drv/usoc  
32-bit ELF kernel module  
/kernel/drv/sparcv9/usoc  
64-bit ELF kernel module
- ATTRIBUTES** See attributes(5) for descriptions of the following attributes:

| ATTRIBUTE TYPE      | ATTRIBUTE VALUE |
|---------------------|-----------------|
| Architecture        | SPARC           |
| Interface stability | Unknown         |
| Availability        | SUNWusoc        |

- SEE ALSO** fctl(7D), sbus(4), fcp(7D), fp(7D), ssd(7D)  
*Writing Device Drivers*  
*Fibre Channel Physical and Signaling Interface (FC-PH) ANSI X3.230: 1994*  
*Fibre Channel Arbitrated Loop (FC-AL) ANSI X3.272-1996*  
*Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA) NCITS TR-19:1998*  
*Fabric Channel Loop Attachment (FC-FLA), NCITS TR-20:1998*
- DIAGNOSTICS** The following messages are logged and may also appear on the system console. On the console these messages are preceded by:
- ```
usoc%d:  
  
where  
usoc%d:  
is the per-port instance number of the usoc controller.  
  
Fibre Channel is ONLINE  
  
The Fibre Channel loop is now online.  
  
Fibre Channel Loop is ONLINE  
  
The Fibre Channel loop is now online.  
  
Fibre Channel Loop is OFFLINE
```

usoc(7D)

The Fibre Channel loop is now offline.

attach failed: device in slave-only slot.

Move soc+ card to another slot.

attach failed: alloc soft state.

Driver did not attach, devices will be inaccessible.

attach failed: bad soft state.

Driver did not attach, devices will be inaccessible.

attach failed: unable to map eeprom

Driver was unable to map device memory; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to map XRAM

Driver was unable to map device memory; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to map registers

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to access status register

Driver was unable to map device registers; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to install interrupt handler

Driver was not able to add the interrupt routine to the kernel. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to access host adapter XRAM

Driver was unable to access device RAM; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: unable to write host adapter XRAM

Driver was unable to write device RAM; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

attach failed: read/write mismatch in XRAM

Driver was unable to verify device RAM; check for bad hardware. Driver did not attach to device, devices will be inaccessible.

NAME	visual_io – Solaris VISUAL I/O control operations
SYNOPSIS	<pre>#include <sys/visual_io.h></pre>
DESCRIPTION	<p>The Solaris VISUAL environment defines a small set of <code>ioctl()</code>s for controlling graphics and imaging devices.</p> <p>One <code>ioctl()</code>, <code>VIS_GETIDENTIFIER</code>, is mandatory, and must be implemented in device drivers for graphics devices using the Solaris VISUAL environment. The <code>VIS_GETIDENTIFIER</code> <code>ioctl()</code> is defined to return a device identifier from the device driver. This identifier must be a uniquely-defined string.</p> <p>Two other sets of <code>ioctl()</code>s exist. One set supports mouse tracking via hardware cursor operations. These are optional, but if a graphics device has hardware cursor support and implements these <code>ioctl()</code>s the mouse tracking performance will be improved.</p> <p>The other set supports the device being the system console device. These are optional, but if a graphics device is to be used as the system console device, it must implement these <code>ioctl()</code>s.</p>
IOCTLS	<p><code>VIS_GETIDENTIFIER</code></p> <p>This <code>ioctl()</code> returns an identifier string to uniquely identify a device used in the Solaris VISUAL environment. This is a mandatory <code>ioctl()</code> and must return a unique string. We suggest that the name be formed as <code><companysymbol><devicetype></code>. For example, the <code>cgsix</code> driver returns <code>SUNWcg6</code>.</p> <p><code>VIS_GETIDENTIFIER</code> takes a <code>vis_identifier</code> structure as its parameter. This structure has the form:</p> <pre>#define VIS_MAXNAMELEN 128 struct vis_identifier { char name[VIS_MAXNAMELEN]; };</pre> <p><code>VIS_GETCURSOR</code> <code>VIS_SETCURSOR</code></p> <p>These <code>ioctl()</code>s fetch and set various cursor attributes, using the <code>vis_cursor</code> structure.</p> <pre>struct vis_cursorpos { short x; /* cursor x coordinate */ short y; /* cursor y coordinate */ }; struct vis_cursorcmap { int version; /* version */ int reserved; unsigned char *red; /* red color map elements */ unsigned char *green; /* green color map elements */ unsigned char *blue; /* blue color map elements */ };</pre>

visual_io(7I)

```
#define VIS_CURSOR_SETCURSOR    0x01 /* set cursor */
#define VIS_CURSOR_SETPOSITION 0x02 /* set cursor position */
#define VIS_CURSOR_SETHOTSPOT  0x04 /* set cursor hot spot */
#define VIS_CURSOR_SETCOLOMAP  0x08 /* set cursor colormap */
#define VIS_CURSOR_SETSHAPE    0x10 /* set cursor shape */
#define VIS_CURSOR_SETALL      \
    (VIS_CURSOR_SETCURSOR | VIS_CURSOR_SETPOSITION | \
     VIS_CURSOR_SETHOTSPOT | VIS_CURSOR_SETCOLOMAP | \
     VIS_CURSOR_SETSHAPE)

struct vis_cursor {
    short set; /* what to set */
    short enable; /* cursor on/off */
    struct vis_cursorpos pos; /* cursor position */
    struct vis_cursorpos hot; /* cursor hot spot */
    struct vis_cursorcmap cmap; /* color map info */
    struct vis_cursorpos size; /* cursor bitmap size */
    char *image; /* cursor image bits */
    char *mask; /* cursor mask bits */
};
```

The `vis_cursorcmap` structure should contain pointers to two elements, specifying the red, green, and blue values for foreground and background.

`VIS_SETCURSORPOS`

`VIS_MOVECURSOR`

These `ioctl()`s fetch and move the current cursor position, using the `vis_cursorpos` structure.

Console optional ioctls

The following set of `ioctl()`s are used by graphics drivers that are to be part of the system console device. All of the `ioctl()`s must be implemented to be a console device. In addition, if the system does not have a prom or the prom goes away during boot, the special standalone `ioctl()`s (listed below) must also be implemented.

The coordinate system for the console device places 0,0 at the upper left corner of the device, with rows increasing toward the bottom of the device and columns increasing from left to right.

`VIS_PUTCMAP`

`VIS_GETCMAP`

Set or get color map entries.

The argument is a pointer to a `vis_cmap` structure, which contains the following fields:

```
struct vis_cmap {
    int    index;
    int    count;
    uchar_t *red;
    uchar_t *green;
    uchar_t *blue;
}
```

`index` is the starting index in the color map where you want to start setting or getting color map entries.

`count` is the number of color map entries to set or get. It also is the size of the red, green, and blue color arrays.

`*red`, `*green`, and `*blue` are pointers to unsigned character arrays which contain the color map info to set or where the color map info is placed on a get.

VIS_DEVINIT

Initializes the graphics driver as a console device.

The argument is a pointer to a `vis_devinit` structure. The graphics driver is expected to allocate any local state information needed to be a console device and fill in this structure.

```
struct vis_devinit {
    int version;
    screen_size_t width;
    screen_size_t height;
    screen_size_t linebytes;
    unit_t size;
    int depth;
    short mode;
};
```

`version` is the version of this structure and should be set to `VIS_CONS_REV`.

`width` and `height` are the width and height of the device. If `mode` (see below) is `VIS_TEXT` then `width` and `height` are the number of characters wide and high of the device. If `mode` is `VIS_PIXEL` then `width` and `height` are the number of pixels wide and high of the device.

`linebytes` is the number of bytes per line of the device.

`size` is the total size of the device in pixels.

`depth` is the pixel depth in bits of the device. Currently supported depths are: 1, 4, 8 and 24.

`mode` is the mode of the device. One of `VIS_PIXEL` (data to be displayed is in bitmap format) or `VIS_TEXT` (data to be displayed is in ascii format).

VIS_DEVFINI

Tells the graphics driver that it is no longer the system console device. There is no argument to this `ioctl()`. The driver is expected to free any locally kept state information related to the console.

VIS_CONS_MODE_CHANGE

Tells the graphics driver that the framebuffer resolution has been reset by the user program. The framebuffer is expected to reload any state information that it is keeping.

The argument to this `ioctl()` is private to the user program and the device driver. That is, the user program may wish to directly change the framebuffer mode and then just use this `ioctl()` to notify the graphics driver or it may pass mode change information along to the graphics driver and have it do the mode change.

VIS_CONSCURS

Describes the size and placement of the cursor on the screen. The graphics driver is expected to display or hide the cursor at the indicated position.

The argument is a pointer to a `vis_conscursor` structure which contains the following fields:

```
struct vis_conscursor {
    int          version;
    screen_pos_t row;
    screen_pos_t col;
    screen_size_t width;
    screen_size_t height;
    color_t      fg_color;
    color_t      bg_color;
    short        action;
};
```

`version` is set to `VIS_CURSOR_VERSION` and should be checked by the driver. If the version does not match, the driver should reject this `ioctl()`.

`row` and `col` are the first row and column (upper left corner of the cursor).

`width` and `height` are the width and height of the cursor.

If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_PIXEL`, then `col`, `row`, `width` and `height` are in pixels. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_TEXT`, then `col`, `row`, `width` and `height` are in characters.

`fg_color` and `bg_color` are the foreground and background color map indexes to use when the `action` (see below) is set to `VIS_DISPLAY_CURSOR`.

`action` is whether to display or hide the cursor. It is set to one of: `VIS_HIDE_CURSOR` or `VIS_DISPLAY_CURSOR`.

VIS_CONSDISPLAY

Display data on the graphics device. The graphics driver is expected to display the data contained in the `vis_display` structure at the specified position on the console.

The `vis_display` structure contains the following fields:

```
struct vis_display {
    int          version;
    screen_pos_t row;
    screen_pos_t col;
    screen_size_t width;
    screen_size_t height;
};
```



```

    uchar_t      *data;
    color_t      fg_color;
    color_t      bg_color;
};

```

version is set to `VIS_DISPLAY_VERSION` and should be checked by the driver. If the version does not match, the driver should reject this `ioctl()`.

`row` and `col` specify the starting row and column to display the data at. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_TEXT`, `row` and `col` are defined to be a character offset from the starting position of the console device. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_PIXEL`, `row` and `col` are defined to be a pixel offset from the starting position of the console device.

`width` and `height` specify the size of the data to be displayed. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_TEXT`, `width` and `height` define the size of data as a rectangle that is `width` characters wide and `height` characters high. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_PIXEL`, `width` and `height` define the size of data as a rectangle that is `width` pixels wide and `height` pixels high.

`*data` is a pointer to the data to be displayed on the console device. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_TEXT`, `data` is an array of ASCII characters to be displayed on the console device. The driver must break these characters up appropriately and display it in the rectangle defined by `row`, `col`, `width`, and `height`. If mode in the `VIS_DEVINIT ioctl()` was set to `VIS_PIXEL`, `data` is an array of bitmap data to be displayed on the console device. The driver must break this data up appropriately and display it in the rectangle defined by `row`, `col`, `width`, and `height`.

The `fg_color` and `bg_color` fields define the foreground and background color map indexes to use when displaying the data. `fg_color` is used for "on" pixels and `bg_color` is used for "off" pixels.

VIS_CONSCOPY

Copy data from one location on the device to another. The driver is expected to copy the specified data. The source data should not be modified. Any modifications to the source data should be as a side effect of the copy destination overlapping the copy source.

The argument is a pointer to a `vis_copy` structure which contains the following fields:

```

struct vis_copy {
    int          version;
    screen_pos_t s_row;
    screen_pos_t s_col;
    screen_pos_t e_row;
    screen_pos_t e_col;
    screen_pos_t t_row;
    screen_pos_t t_col;
};

```

```

        short      direction;
    };

```

version is set to VIS_COPY_VERSION and should be checked by the driver. If the version does not match, the driver should reject this ioctl().

s_row, s_col, e_row, and e_col define the source rectangle of the copy. s_row and s_col are the upper left corner of the source rectangle. e_row and e_col are the lower right corner of the source rectangle. If mode in the VIS_DEVINIT ioctl() was set to VIS_TEXT, s_row, s_col, e_row, and e_col are defined to be character offsets from the starting position of the console device. If mode in the VIS_DEVINIT ioctl() was set to VIS_PIXEL, s_row, s_col, e_row, and e_col are defined to be pixel offsets from the starting position of the console device.

t_row and t_col define the upper left corner of the destination rectangle of the copy. The entire rectangle is copied to this location. If mode in the VIS_DEVINIT ioctl() was set to VIS_TEXT, t_row, and t_col are defined to be character offsets from the starting position of the console device. If mode in the VIS_DEVINIT ioctl() was set to VIS_PIXEL, t_row, and t_col are defined to be pixel offsets from the starting position of the console device.

direction specifies which way to do the copy. If direction is VIS_COPY_FORWARD the graphics driver should copy data from position (s_row, s_col) in the source rectangle to position (t_row, t_col) in the destination rectangle. If direction is VIS_COPY_BACKWARDS the graphics driver should copy data from position (e_row, e_col) in the source rectangle to position (t_row+(e_row-s_row), t_col+(e_col-s_col)), in the destination rectangle.

The next set of console ioctl()s are used on systems which don't have a prom. Normally, standalones use the system prom to display characters on the system console device. On systems without a prom, standalones use the kernel drivers to display characters on the system console device. When implementing these ioctl()s, you can not use any of the locking primitives or the copy routines from the DDI. Furthermore other DDI services may or may not work and should be avoided.

VIS_STAND_CONSCURSOR

Should perform the same tasks as VIS_CONSCURSOR except that it must follow the above restrictions. It takes in as an argument a vis_cursor structure.

VIS_STAND_CONSDISPLAY

Should perform the same tasks as VIS_CONSDISPLAY except that it must follow the above restrictions. It takes in as an argument a vis_display structure.

VIS_STAND_CONSCOPY

Should perform the same tasks as VIS_CONSCOPY except that it must follow the above restrictions. It takes in as an argument a vis_copy structure.

SEE ALSO ioctl(2)

NAME	vni – STREAMS virtual network interface driver
DESCRIPTION	<p>The <code>vni</code> pseudo device is a multi-threaded, loadable, clonable, STREAMS pseudo-device supporting the connectionless Data Link Provider Interface <code>dlpi(7P)</code> Style 2. Note that DLPI is intended to interact with IP, meaning that DLPI access to applications is not supported. (For example, <code>snoop</code> fails on the <code>vni</code> interface.)</p> <p>The <code>vni</code> device is a software-only interface and does not send or receive data. The device provides a DLPI upper interface that identifies itself to IP with a private media type. It can be configured via <code>ifconfig(1M)</code> and can have IP addresses assigned to it, making aliases possible.</p> <p>The <code>vni</code> pseudo device is particularly useful in hosting an IP address when used in conjunction with the <code>'usesrc'</code> <code>ifconfig</code> option (see <code>ifconfig(1M)</code> for examples). The logical instances of the device can also be used to host addresses as an alternative to hosting them over the loopback interface.</p> <p>Multicast is not supported on this device. More specifically, the following options return an error when used with an address specified on <code>vni</code>: <code>IP_MULTICAST_IF</code>, <code>IP_ADD_MEMBERSHIP</code>, <code>IP_DROP_MEMBERSHIP</code>, <code>IPV6_MULTICAST_IF</code>, <code>IPV6_JOIN_GROUP</code>, <code>IPV6_LEAVE_GROUP</code>. In addition, broadcast is not supported.</p> <p>Because there is no physical hardware configured below it, no traffic can be received through nor transmitted on a virtual interface. All packet transmission and reception is accomplished with existing physical interfaces and tunnels. Because applications that deal with packet transmission and reception (such as packet filters) cannot filter traffic on virtual interfaces, you cannot set up a packet filter on a virtual interface. Instead, you should configure the policy rules to apply to the physical interfaces and tunnels, and if necessary, use the virtual IP addresses themselves as part of the rule configuration. Also, note that the virtual interface cannot be part of an IP multipathing (IPMP) group.</p>
FILES	<p><code>/dev/vni</code> 64-bit ELF kernel driver</p>
SEE ALSO	<code>ifconfig(1M)</code> , <code>in.mpathd(1M)</code> , <code>ip(7P)</code> , <code>ip6(7P)</code>

volfs(7FS)

NAME volfs – Volume Management file system

DESCRIPTION `volfs` is the Volume Management file system rooted at `root_dir`. The default location for `root_dir` is `/vol`, but this can be overridden using the `-d` option of `vold` (see `vold(1M)`). This file system is maintained by the Volume Management daemon, `vold`, and will be considered to be `/vol` for this description.

Media can be accessed in a logical manner (no association with a particular piece of hardware), or a physical manner (associated with a particular piece of hardware).

Logical names for media are referred to through `/vol/dsk` and `/vol/rdsk`. `/vol/dsk` provides block access to random access devices. `/vol/rdsk` provides character access to random access devices.

The `/vol/rdsk` and `/vol/dsk` directories are mirrors of one another. Any change to one is reflected in the other immediately. The `dev_t` for a volume will be the same for both the block and character device.

The default permissions for `/vol` are `mode=0555, owner=root, group=sys`. The default permissions for `/vol/dsk` and `/vol/rdsk` are `mode=01777, owner=root, group=sys`.

Physical references to media are obtained through `/vol/dev`. This hierarchy reflects the structure of the `/dev` name space. The default permissions for all directories in the `/vol/dev` hierarchy are `mode=0555, owner=root, group=sys`.

`mkdir(2)`, `rmdir(2)`, `unlink(2)` (`rm`), `symlink(2)` (`ln -s`), `link(2)` (`ln`), and `rename(2)` (`mv`) are supported, subject to normal file and directory permissions.

The following system calls are not supported in the `/vol` filesystem: `creat(2)`, only when creating a file, and `mknod(2)`.

If the media does not contain file systems that can be automatically mounted by `rmmount(1M)`, users can gain access to the media through the following `/vol` locations:

Location	State of Media
<code>/vol/dev/diskette0/unnamed_floppy</code>	formatted unnamed floppy-block device access
<code>/vol/dev/rdiskette0/unnamed_floppy</code>	formatted unnamed floppy-raw device access
<code>/vol/dev/diskette0/unlabeled</code>	unlabeled floppy-block device access
<code>/vol/dev/rdiskette0/unlabeled</code>	unlabeled floppy-raw device access
<code>/vol/dev/dsk/c0t6/unnamed_cdrom</code>	CD-ROM-block device access

Location	State of Media
/vol/dev/rdisk/c0t6/unnamed_cdrom	CD-ROM-raw device access

For more information on the location of CD-ROM and floppy media, see *System Administration Guide: Basic Administration* or `rmmount(1M)`.

Partitions

Some media support the concept of a partition. If the label identifies partitions on the media, the name of the media becomes a directory with partitions under it. Only valid partitions are represented. Partitions cannot be moved out of a directory.

For example, if disk volume 'foo' has three valid partitions, 0, 2, and 5, then:

```
/vol/dsk/foo/s0
/vol/dsk/foo/s2
/vol/dsk/foo/s5
```

for block access and

```
/vol/rdsk/foo/s0
/vol/rdsk/foo/s2
/vol/rdsk/foo/s5
```

for character access.

If a volume is relabeled to reflect different partitions, the name space changes to reflect the new partition layout.

A format program can check to see if there are others with the volume open and not allow the format to occur if it is. Volume Management, however, does not explicitly prevent the rewriting of a label while others have the volume open. If a partition of a volume is open, and the volume is relabeled to remove that partition, it will appear exactly as if the volume were missing. A notify event will be generated and the user may cancel the operation with `volcancel(1)`, if desired.

SEE ALSO

`volcancel(1)`, `volcheck(1)`, `volmissing(1)`, `rmmount(1M)`, `vold(1M)`, `rmmount.conf(4)`, `vold.conf(4)`

System Administration Guide: Basic Administration

vuidmice(7M)

NAME	vuidmice, vuidm3p, vuidm4p, vuidm5p, vuid2ps2, vuid3ps2 – converts mouse protocol to Firm Events				
SYNOPSIS	<pre>#include <sys/stream.h> #include <sys/vuid_event.h> #include <sys/vuid_wheel.h> int ioctl(fd, I_PUSH, vuidm3p); int ioctl(fd, I_PUSH, vuidm4p); int ioctl(fd, I_PUSH, vuidm5p); int ioctl(fd, I_PUSH, vuid2ps2); int ioctl(fd, I_PUSH, vuid3ps2);</pre>				
DESCRIPTION	<p>The STREAMS modules vuidm3p, vuidm4p, vuidm5p, vuid2ps2, and vuid3ps2 convert mouse protocols to Firm events. The Firm event structure is described in <sys/vuid_event.h>. Pushing a STREAMS module does not automatically enable mouse protocol conversion to Firm events. The STREAMS module state is initially set to raw or VUID_NATIVE mode which performs no message processing. You must change the state to VUID_FIRM_EVENT mode to initiate mouse protocol conversion to Firm events. This can be accomplished by the following code:</p> <pre>int format; format = VUID_FIRM_EVENT; ioctl(fd, VUIDSFORMAT, &format);</pre> <p>You can also query the state of the STREAMS module by using the VUIDGFORMAT option.</p> <pre>int format; int fd; /* file descriptor */ ioctl(fd, VUIDGFORMAT, &format); if (format == VUID_NATIVE); /* The state of the module is in raw mode. * Message processing is not enabled. */ if (format == VUID_FIRM_EVENT); /* Message processing is enabled. * Mouse protocol conversion to Firm events * are performed.</pre> <p>The remainder of this section describes the processing of STREAMS messages on the read- and write-side.</p>				
Read Side Behavior	<table><tr><td>M_DATA</td><td>Incoming messages are queued and converted to Firm events.</td></tr><tr><td>M_FLUSH</td><td>The read queue of the module is flushed of all its data messages and all data in the record being accumulated are also flushed. The message is passed upstream.</td></tr></table>	M_DATA	Incoming messages are queued and converted to Firm events.	M_FLUSH	The read queue of the module is flushed of all its data messages and all data in the record being accumulated are also flushed. The message is passed upstream.
M_DATA	Incoming messages are queued and converted to Firm events.				
M_FLUSH	The read queue of the module is flushed of all its data messages and all data in the record being accumulated are also flushed. The message is passed upstream.				

Write Side Behavior

M_IOCTL	Messages sent downstream as a result of an <code>ioctl(2)</code> system call. The two valid <code>ioctl</code> options processed by the <code>vuidmice</code> modules are <code>VUIDGFORMAT</code> and <code>VUIDSFORMAT</code> .
M_FLUSH	The write queue of the module is flushed of all its data messages and the message is passed downstream.
VUIDGFORMAT	This option returns the current state of the STREAMS module. The state of the <code>vuidmice</code> STREAMS module may either be <code>VUID_NATIVE</code> (no message processing) or <code>VUID_FIRM_EVENT</code> (convert to Firm events).
VUIDSFORMAT	This option sets the state of the STREAMS module to <code>VUID_FIRM_EVENT</code> . If the state of the STREAMS module is already in <code>VUID_FIRM_EVENT</code> , this option is non-operational. It is not possible to set the state back to <code>VUID_NATIVE</code> once the state becomes <code>VUID_FIRM_EVENT</code> . To disable message processing, pop the STREAMS module out by calling <code>ioctl(fd, lI_POP, vuid*)</code> .

The following wheel support `ioctl`s are defined for PS/2 mouse only:

VUIDGWHEELCOUNT	This <code>ioctl</code> takes a pointer to an integer as argument and sets the value of the integer to the number of wheels available on this device.
VUIDGWHEELINFO	This command returns static information about the wheel that does not change while a device is in use. Currently the only information defined is the wheel orientation which is either <code>VUID_WHEEL_FORMAT_VERTICAL</code> or <code>VUID_WHEEL_FORMAT_HORIZONTAL</code> . <pre> typedef struct { int vers; int id; int format; } wheel_info; </pre> <p>The <code>ioctl</code> takes a pointer to "wheel_info" structure with the "vers" set to the current version of the "wheel_info" structure and "id" set to the id of the wheel for which the information is desired.</p>
VUIDSWHEELSTATE VUIDGWHEELSTATE	<code>VUIDSWHEELSTATE</code> sets the state of the wheel to that specified in the <code>stateflags</code> . <code>VUIDGWHEELSTATE</code> returns the current state settings in the <code>stateflags</code> field.

stateflags is an OR'ed set of flag bits. The only flag currently defined is VUID_WHEEL_STATE_ENABLED.

When stateflags is set to VUID_WHEEL_STATE_ENABLED the module converts motion of the specified wheel into VUID events and sends those up stream.

Wheel events are disabled by default.

Applications that want to change a flag should first get the current flags and then change only the bit they want.

```
typedef struct {
    int      vers;
    int      id;
    uint32_t stateflags;
} wheel_state;
```

These ioctls take pointer to 'wheel_state' as an argument with the 'vers' and 'id' members filled up. These members have the same meaning as that for 'VUIDGWHEELINFO' ioctl.

Mouse Configurations

Module	Protocol Type	Device
vuidm3p	3-Byte Protocol Microsoft 2 Button Serial Mouse	/dev/tty*
vuidm4p	4-Byte Protocol Logitech 3 Button Mouseman	/dev/tty*
vuidm5p	Logitech 3 Button Bus Mouse Microsoft Bus Mouse	/dev/logi/ dev /msm
vuid2ps2	PS/2 Protocol 2 Button PS/2 Compatible Mouse	/dev/kdmouse
vuid3ps2	PS/2 Protocol 3 Button PS/2 Compatible Mouse	/dev/kdmouse

ATTRIBUTES

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO

attributes(5)

wrsm(7D)

NAME	wrsm – WCI Remote Shared Memory (WRSM) device driver
SYNOPSIS	wci@<slot>,0:wrsm wrsm@<instance>:ctrl wrsm@ffff,0:admin
DESCRIPTION	<p>The wrsm driver is a nexus driver that manages Sun Fire Link devices and wrsm controllers.</p> <p>A WCI device on a Sun Fire Link board is attached directly to the host system bus and provides clustering communication between Solaris instances that are memory transaction-based. The WCI acts as a memory controller on the system backplane. The wrsm driver programs registers on the WCI to accept network read/write requests on certain exported cluster addresses from incoming links. The registers translate the requests into local read/write bus transactions that use local physical memory ranges that you specify. The driver programs additional WCI registers to forward local system backplane read/write transactions within a particular physical address range to a remote WCI. A WCI device in the format wci@slot,0:wrsm appears in the device tree.</p> <p>A wrsm controller is a pseudo device that manages a set of WCIs. A device entry in the format wrsm@<instance>:ctrl appears in the device tree. A wrsm controller presents a Sun proprietary protocol to clients, enabling them to set up the network and to communicate through the WCIs. To configure a wrsm controller, you download a configuration into the driver using the wrsmconf(1M) command or through other external WCI network management software. Status information on each WCI and wrsm controller is available by using the wrsmstat(1M) command.</p> <p>The wrsm admin device is used internally by the driver to manage the I/O addresses associated with remote memory. A device entry in the format wrsm@ffff,0:admin appears in the device tree</p>
FILES	/platform/sun4u/kernel/drv/sparcv9/wrsm ELF kernel module
SEE ALSO	wrsmconf(1M), wrsmstat(1M) <i>Writing Device Drivers</i>
DIAGNOSTICS	<p>The messages described below may appear on the system console as well as being logged. These messages generally include the string wrsm%d, where %d is the instance number of the wrsm device. The message context indicates whether the device is a WCI or a wrsm controller. Some messages include the string wci %a, where %a is the bus slot of the WCI device.</p> <p>wrsm%d: unable to map register set %d Driver was unable to map device registers; check for bad hardware. Driver did not attach device, device will be inaccessible.</p> <p>wrsm_detach:cf_remove_controller failed for wrsm%d Driver did not detach device; device is inaccessible.</p>

wrsm_detach:cf_remove_wci failed for wrsm%d

Driver did not detach device. This WCI is the last WCI in wrsm controller.

register_controller of wrsm%d failed with error %d

The wrsm controller could not register with the Sun proprietary protocol framework. Communication is not possible through this controller.

wrsm%d, wci %a, SRAM CE ERROR, at address: 0x%x, syndrome: 0x%x

There was a correctable error in the WCI's SRAM. This indicates that the memory on this WCI module should be replaced.

wrsm%d, wci %a, SRAM UE ERROR, at address: 0x%x, syndrome: 0x%x

There was an uncorrectable error in the WCI's SRAM. This indicates that the memory on this WCI module should be replaced. In addition, attempts to access local memory from remote nodes may fail.

wrsmd(7D)

NAME	wrsmd – WCI Remote Shared Memory (WRSM) DLPI driver
SYNOPSIS	wrsmd
DESCRIPTION	<p>The <code>wrsmd</code> device driver is a pseudo driver that presents a type II DLPI interface. The driver uses a Sun proprietary interface over the <code>wrsm</code> driver to provide IP-based communication over the WCI network.</p> <p>Before using DLPI over a WCI network, you must first configure the <code>wrsm</code> controller on that network. See <code>wrsmconf(1M)</code> for more information. Each <code>wrsmd</code> device is associated with a <code>wrsm</code> controller with a matching id.</p>
FILES	<code>/platform/sun4u/kernel/drv/sparcv9/wrsmd</code> ELF kernel module
SEE ALSO	<code>wrsmconf(1M)</code> , <code>wrsm(7D)</code> <i>Writing Device Drivers</i>

NAME	wscons – workstation console				
SYNOPSIS	<pre>#include <sys/strredir.h> ioctl(<i>fd</i>, SRIOCSREDIR, <i>target</i>); ioctl(<i>fd</i>, SRIOCISREDIR, <i>target</i>);</pre>				
DESCRIPTION	<p>The <code>wscons</code> workstation console consists of a workstation keyboard and frame buffer that act together to emulate an ASCII terminal. It includes a redirection facility that allows I/O issued to the workstation console to be diverted to a STREAMS device, enabling window systems to redirect output that would otherwise appear directly on the frame buffer in corrupted form.</p>				
Redirection	<p>The <code>wscons</code> redirection facility maintains a list of devices that are designated as redirection targets through the <code>SRIOCSREDIR</code> <code>ioctl</code> described below. Only the current entry is active; when the active entry is closed, the most recent remaining entry becomes active. The active entry acts as a proxy for the device being redirected and handles all <code>read(2)</code>, <code>write(2)</code>, <code>ioctl(2)</code>, and <code>poll(2)</code> calls issued against the redirectee.</p> <p>The <code>ioctls</code> described below control the redirection facility. In both cases, <i>fd</i> is a descriptor for the device being redirected (or workstation console) and <i>target</i> is a descriptor for a STREAMS device.</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>SRIOCSREDIR</code></td> <td>Designates <i>target</i> as the source and destination of I/O ostensibly directed to the device denoted by <i>fd</i>.</td> </tr> <tr> <td><code>SRIOCISREDIR</code></td> <td>Returns 1 if <i>target</i> names the device currently acting as proxy for the device denoted by <i>fd</i>, and 0 if it is not.</td> </tr> </table>	<code>SRIOCSREDIR</code>	Designates <i>target</i> as the source and destination of I/O ostensibly directed to the device denoted by <i>fd</i> .	<code>SRIOCISREDIR</code>	Returns 1 if <i>target</i> names the device currently acting as proxy for the device denoted by <i>fd</i> , and 0 if it is not.
<code>SRIOCSREDIR</code>	Designates <i>target</i> as the source and destination of I/O ostensibly directed to the device denoted by <i>fd</i> .				
<code>SRIOCISREDIR</code>	Returns 1 if <i>target</i> names the device currently acting as proxy for the device denoted by <i>fd</i> , and 0 if it is not.				
ANSI Standard Terminal Emulation	<p>On SPARC systems, the PROM monitor emulates an ANSI X3.64 terminal.</p> <p>On x86 systems, the Solaris console subsystem provides ANSI X3.64 emulation.</p> <p>Note: The VT100 adheres the ANSI X3.64 standard. However, because the VT100 features nonstandard extensions to ANSI X3.64, it is incompatible with Sun terminal emulators.</p> <p>The SPARC console displays 34 lines of 80 ASCII characters per line. The x86 console displays 25 lines of 80 ASCII characters per line. Devices with smaller text capacities may display less. On SPARC systems, the <code>screen-#rows</code> and <code>screen-#columns</code> should be set to 34 or 80 respectively, or text capacities will vary from those described above. On SPARC systems, the <code>screen-#rows</code> and <code>screen-#columns</code> fields are stored in NVRAM/EEPROM. See <code>eeprom(1M)</code> for more information. Both SPARC and IA consoles offer scrolling, (x, y) cursor addressing ability, and a number of other control functions.</p>				

wscnons(7D)

The console cursor marks the current line and character position on the screen. ASCII characters between 0x20 (space) and 0x7E (tilde) inclusive are printing characters. When a print character is written to the console (and is not part of an escape sequence), it is displayed at the current cursor position and the cursor moves one position to the right on the current line.

On SPARC based systems, later PROM revisions have the full 8-bit ISO Latin-1 (*ISO 8859-1*) character set. Earlier PROM revisions display characters in the range 0xA0 through 0xFE as spaces.

When the cursor is at the right edge of the screen, it moves to the first character position on the next line. When the cursor is at the screen's right-bottom edge, the line-feed function is performed (see CTRL-J below). The line-feed function scrolls the screen up by one or more lines before moving the cursor to the first character position on the next line.

Control Sequence Syntax

The `wscnons` console defines a number of control sequences that may occur during input. When a control sequence is written to the console, it affects one of the control functions described below. Control sequences are not displayed on screen.

A number of control sequences (or control character functions) are of the form:

CTRL-*x*

where *x* represents a single character, such as CTRL-J for a line feed.

Other ANSI control sequences are of the form:

ESC [*params char*

Note – Spaces are included only for readability; these characters must occur in the given sequence without the intervening spaces.

ESC	ASCII escape character (ESC, CTRL-[, 0x1B).
[Left square bracket '[' (0x5B).
<i>params</i>	Sequence of zero or more decimal numbers made up of digits between 0 and 9, separated by semicolons. Parameters are represented by <i>n</i> in the syntax descriptions for escape sequence functions.
<i>char</i>	Function character, which is different for each control sequence and it represented by <i>x</i> in the syntax descriptions for control character functions.

In the following examples of syntactically valid escape sequences, ESC represent the single ASCII character, Escape:

ESC[m Select graphic rendition with default parameter

ESC[7m	Select graphic rendition with reverse image
ESC[33;54H	Set cursor position
ESC[123;456;0;;3;B	Move cursor down

Syntactically valid control characters and ANSI escape sequences that are not currently interpreted by the console are ignored.

Each control function requires a specified number of parameters. If fewer parameters are supplied, the remaining parameters (which certain exceptions that are noted below) default to 1. For example, if more than the required number of parameters are supplied, only the last n is used, where n is the number required by that particular command character.

Parameters which are omitted or set to 0 are reset to the default value of 1 (with certain exceptions). For example, the command character M requires one parameter. $ESC [; M$, $ESC [0M$, $ESC [M$ and $ESC [23 ; 15 ; 32 ; 1M$ are all equivalent to $ESC [1M$ and provide a parameter value of 1. Note that $ESC [; 5M$ (interpreted as $'ESC [5M'$) is *not* equivalent to $ESC [5 ; M$ (interpreted as $'ESC [5 ; 1M'$) which is ultimately interpreted as $'ESC [1M'$).

ANSI Control Functions

The following paragraphs specify the ANSI control functions implemented by the console. Each description provides:

- Control sequence syntax
- Hexadecimal equivalent of control characters where applicable
- Control function name and ANSI or Sun abbreviation (if any).
- Description of parameters required, if any
- Description of the control function
- Initial setting of the mode for functions that set a mode. To restore the initial settings, use the SUNRESET escape sequence.

Control Character Functions

The `wscons` control character functions are:

Bell (BEL),
CTRL-G
0x7

Used for consoles that are not equipped with an audible bell. Current Sun workstation models also flash the screen if the keyboard is not the console input device.

Backspace (BS),
CTRL-H,
0x8

The cursor moves one position to the left on the current line. If it is already at the left edge of the screen, no change takes place.

Tab (TAB),
CTRL-I,

0x9

The cursor moves right on the current line to the next tab stop. The tab stops are fixed at every multiple of eight columns. If the cursor is already at the right edge of the screen, nothing change takes place. Otherwise, the cursor moves right a minimum of one and a maximum of eight character positions.

Line-feed (LF),

CTRL-J,

0xA

The cursor, while remaining at the same character position on the line, moves down one line. If the cursor is at the bottom line, the screen either scrolls up or wraps around depending on the setting of an internal variable n (initially 1) . The internal variable can be changed using the ESC [r control sequence. If n is greater than zero, the entire screen (including the cursor) is scrolled up by n lines before executing the line-feed. The top n lines scroll off the screen and are lost. New blank lines n scroll onto the bottom of the screen. After scrolling, move the cursor down one line to execute the line feed.

If n is zero, wrap-around mode is entered. The ESC [1 r exits back to scroll mode. If a line-feed occurs on the bottom line in wrap mode, the cursor goes to the same character position in the top line of the screen. During line-feeds, the line that the cursor moves to is cleared and no scrolling occurs. Wrap-around mode is not implemented in the window system.

On SPARC based systems, the speed at which the screen scrolls is dependent on the amount of data waiting to be printed. Whenever a scroll occurs and the console is in normal scroll mode (ESC [1 r), it scans the rest of the data awaiting printing to see how many line-feeds occur in it. This scan stops when the console finds a control character from the set {VT, FF, SO, SI, DLE, DC1, DC2, DC3, DC4, NAK, SYN, ETB, CAN, EM, SUB, ESC, FS, GS, RS, US} . At that point, the screen is scrolled by n lines ($n \geq 1$) and processing continues. The scanned text is processed normally and fills in the newly created lines. As long as escape codes or other control characters are not intermixed with the text, this results in faster scrolling

Reverse Line-feed,

CTRL-K,

0xB

While remaining at the same character position on the line, the cursor moves up one line. If the cursor is already at the top line, no change takes place.

Form-feed (FF)

CTRL-L,

0xC

The cursor is positioned to the home position (upper-left corner) and the entire screen is cleared.

Return (CR),

CTRL-M,

0xD

The cursor moves to the leftmost character position on the current line.

Escape Sequence Functions

The `wscons` escape sequence functions are:

Escape (ESC),
CTRL-[,
0x1B

The escape character. Escape initiates a multi-character control sequence.

Insert Character (ICH)
ESC[#@

Takes one parameter, n (default 1). Inserts n spaces at the current cursor position. The current line, starting at the current cursor position inclusive, is shifted to the right by n character positions to make room for the spaces. The rightmost n character positions shift off the line and are lost. The position of the cursor is unchanged.

Cursor Up (CUU),
ESC[#A

Takes one parameter, n (default 1). Moves the cursor up n lines. If the cursor is fewer than n lines from the top of the screen, moves the cursor to the topmost line on the screen. The character position of the cursor on the line is unchanged.

Cursor Down (CUD),
ESC[#B

Takes one parameter, (default 1). Moves the cursor down n lines. If the cursor is fewer than n lines from the bottom of the screen, move the cursor to the last line on the screen. The character position of the cursor on the line is unchanged.

Cursor Forward (CUF),
ESC[#C

Takes one parameter, n (default 1). Moves the cursor to the right by n character positions on the current line. If the cursor is fewer than n positions from the right edge of the screen, moves the cursor to the rightmost position on the current line.

Cursor Backward (CUB),
ESC[#D

Takes one parameter, n (default 1). Moves the cursor to the left by n character positions on the current line. If the cursor is fewer than n positions from the left edge of the screen, moves the cursor to the leftmost position on the current line.

Cursor Next Line (CNL),
ESC[#E

Takes one parameter, n (default 1). Positions the cursor at the leftmost character position on the n -th line below the current line. If the current line is less than n lines from the bottom of the screen, positions the cursor at the leftmost character position on the bottom line.

Horizontal and Vertical Position (HVP),
ESC[#1;#2f

or

Cursor Position (CUP),
ESC[#1;#2H

Takes two parameters, $n1$ and $n2$ (default 1, 1). Moves the cursor to the $n2$ -th character position on the $n1$ -th line. Character positions are numbered from 1 at the left edge of the screen; line positions are numbered from 1 at the top of the screen. Hence, if both parameters are omitted, the default action moves the cursor to the home position (upper left corner). If only one parameter is supplied, the cursor moves to column 1 of the specified line.

Erase in Display (ED),
ESC[J

Takes no parameters. Erases from the current cursor position inclusive to the end of the screen, that is, to the end of the current line and all lines below the current line. The cursor position is unchanged.

Erase in Line (EL),
ESC[K

Takes no parameters. Erases from the current cursor position inclusive to the end of the current line. The cursor position is unchanged.

Insert Line (IL),
ESC[#L

Takes one parameter, n (default 1). Makes room for n new lines starting at the current line by scrolling down by n lines the portion of the screen from the current line inclusive to the bottom. The n new lines at the cursor are filled with spaces; the bottom n lines shift off the bottom of the screen and are lost. The position of the cursor on the screen is unchanged.

Delete Line (DL),
ESC[#M

Takes one parameter, n (default 1). Deletes n lines beginning with the current line. The portion of the screen from the current line inclusive to the bottom is scrolled upward by n lines. The n new lines scrolling onto the bottom of the screen are filled with spaces; the n old lines beginning at the cursor line are deleted. The position of the cursor on the screen is unchanged.

Delete Character (DCH),
ESC[#P

Takes one parameter, n (default 1). Deletes n characters starting with the current cursor position. Shifts the tail of the current line to the left by n character positions from the current cursor position, inclusive, to the end of the line. Blanks are shifted into the rightmost n character positions. The position of the cursor on the screen is unchanged.

Select Graphic Rendition (SGR),
ESC[#m

Takes one parameter, n (default 0). Note that unlike most escape sequences, the parameter defaults to zero if omitted. Invokes the graphic rendition specified by the parameter. All following printing characters in the data stream are rendered according to the parameter until the next occurrence of this escape sequence in the data stream. Currently only two graphic renditions are defined:

0 Normal rendition

7 Negative (reverse) image

Negative image displays characters as white-on-black if the screen mode is currently black-on white, and vice-versa. Any non-zero value of *n* is currently equivalent to 7 and selects the negative image rendition.

On x86 systems only, the following ISO 6429-1983 graphic rendition values support color text:

30	black foreground
31	red foreground
32	green foreground
33	brown foreground
34	blue foreground
35	magenta foreground
36	cyan foreground
37	white foreground
40	black background
41	red background
42	green background
43	brown background
44	blue background
45	magenta background
46	cyan background
47	white background

Black On White (SUNBOW),

ESC[p

Takes no parameters. Sets the screen mode to black-on-white. If the screen mode is already black-on-white, has no effect. In this mode spaces display as solid white, other characters as black-on-white. The cursor is a solid black block. Characters displayed in negative image rendition (see 'Select Graphic Rendition' above) is white-on-black. This comprises the initial setting of the screen mode on reset.

White On Black (SUNWOB),

ESC[q

Takes no parameters. Sets the screen mode to white-on-black. If the screen mode is already white-on-black, has no effect. In this mode spaces display as solid black, other characters as white-on-black. The cursor is a solid white block. Characters displayed in negative image rendition (see 'Select Graphic Rendition' above) is black-on-white in this mode. The initial setting of the screen mode on reset is black on white.

ESC[#r

Set Scrolling (SUNSCRL)

Takes one parameter, *n* (default 0). Sets to *n* an internal register which determines how many lines the screen scrolls up when a line-feed function is performed with the cursor on the bottom line. A parameter of 2 or 3 introduces a small amount of jump when a scroll occurs. A parameter of 34 clears the screen rather than scrolling. The initial setting is 1 on reset.

A parameter of zero initiates wrap mode instead of scrolling. If a linefeed occurs on the bottom line during wrap mode, the cursor goes to the same character position in the top line of the screen. When a line feed occurs, the line that the cursor moves to is cleared and no scrolling occurs. ESC [1 r exits back to scroll mode.

For more information, see the description of the Line-feed (CTRL-J) control function above.

ESC[s

Reset terminal emulator (SUNRESET)

Takes no parameters. Resets all modes to default, restores current font from PROM. Screen and cursor position are unchanged.

RETURN VALUES

When there are no errors, the redirection ioctls have return values as described above. Otherwise, they return -1 and set `errno` to indicate the error. If the target stream is in an error state, `errno` is set accordingly.

If the *target* stream is in an error state, `errno` is set accordingly.

ERRORS

EBADF *target* does not denote an open file.

ENOSTR *target* does not denote a STREAMS device.

FILES

/dev/wscons Workstation console, accessed via the redirection facility

/dev/systty Devices that must be opened for the SRIOCSREDIR and SRIOCISREDIR ioctls.

/dev/syscon Access system console

/dev/console Access system console

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Stable

SEE ALSO

`cvcd(1M)`, `eeprom(1M)`, `ioctl(2)`, `poll(2)`, `read(2)`, `write(2)`, `cvc(7D)`, `console(7D)`

WARNINGS

The redirection ioctls block while there is I/O outstanding on the device instance being redirected. If you try to redirect the workstation console while there is a outstanding read, the workstation console will hang until the read completes.

NOTES | The `cvc` facility supersedes the SunOS `wscons` facility and should not be used with `wscons`.

xge(7D)

NAME	xge – S2IO Xframe 10Gigabit Ethernet Network Adapter driver
SYNOPSIS	<code>/dev/xge</code>
DESCRIPTION	<p>The <code>xge</code> 10 Gigabit Ethernet driver is a multi-threaded, loadable, clonable, GLD-based STREAMS driver supporting the Data Link Provider Interface, <code>dlpi(7P)</code>, on S2IO Xframe 10-Gigabit Ethernet Network Adapter.</p> <p>The <code>xge</code> driver functions includes controller initialization, frame transmit and receive, promiscuous and multicast support, TCP and UDP checksum offload (IPv4 and IPv6), 9622-byte jumbo frame, and error recovery and reporting.</p> <p>The <code>xge</code> driver and hardware support the 10GBase-SR/W, LR/W, and ER/W 802.3 physical layer.</p>
APPLICATION PROGRAMMING INTERFACE	<p>The cloning, character-special device <code>/dev/xge</code> is used to access all Xframe devices installed within the system.</p> <p>The <code>xge</code> driver is dependent on <code>/kernel/misc/gld</code>, a loadable kernel module that provides the <code>xge</code> driver with the DLPI and STREAMS functionality required of a LAN driver. See <code>gld(7D)</code> for more details on supported primitives.</p> <p>The values returned by the driver in the <code>DL_INFO_ACK</code> primitive in response to the <code>DL_INFO_REQ</code> are as follows:</p> <ul style="list-style-type: none">■ Maximum SDU is 9600.■ Minimum SDU is 0.■ DSLAP address length is 8 bytes.■ MAC type is <code>DL_ETHER</code>.■ SAP length value is -2 meaning the physical address component is followed immediately by a 2-byte sap component within the DLSAP address.■ Broadcast address value is Ethernet/IEEE broadcast address (FF:FF:FF:FF:FF:FF).
CONFIGURATION	<p>By default, the <code>xge</code> driver works without any configuration file.</p> <p>You can check the running-time status of a device instance using <code>ndd(1M)</code>. Currently, the driver provides an interface to print all hardware statistics.</p> <p>For example, to print statistics of device <code>xge0</code>:</p> <pre>#ndd /dev/xge0 stats tmac_data_octets 772 tmac_frms 15 tmac_drop_frms 0 tmac_bcst_frms 6 tmac_mcst_frms 6 ... rmac_vld_frms 13 rmac_fcs_err_frms 0</pre>

```

rmac_drop_frms 0
rmac_vld_bcst_frms 7
rmac_vld_mcst_frms 11
rmac_out_rng_len_err_frms 0
rmac_in_rng_len_err_frms 0
rmac_long_frms 0

```

...

```

not_traffic_intr_cnt 242673
traffic_intr_cnt 28

```

...

FILES /dev/xge xge special character device

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO ndd(1M), attributes(5), gld(7D), streamio(7I), dlpi(7P)

Writing Device Drivers

STREAMS Programming Guide

Network Interfaces Programmer's Guide

xmemfs(7FS)

NAME	xmemfs – extended memory file system
SYNOPSIS	<pre>#include <sys/mount.h> mount(special_file, directory, MS_DATA, "xmemfs", dataptr, datalen);</pre>
DESCRIPTION	<p>The xmemfs file system is an extended memory file system that provides an efficient mechanism for managing and accessing physical memory that exceeds 4 Gbytes in size. Currently, the xmemfs file system is supported on IA32 architecture systems only.</p> <p>The Physical Address Extension (PAE) is the xmemfs internal processor feature that enables a 36-bit physical memory address that supports up to 64 Gbytes of physical memory. Once mounted, the xmemfs file system provides standard file operations and semantics on directories and regular files only. Because xmemfs does not allow execute permissions to be set on regular files, execution of object files is prevented.</p> <p>With xmemfs, the <code>special_file</code> argument, (typically the device on which file systems reside), is ignored and serves only as a placeholder. File data and metadata in xmemfs are always memory-resident. The <code>dataptr</code> argument must (at a minimum) contain the required <code>size</code> specific option. See <code>mount_xmemfs(1M)</code> for more information.</p> <p>Because xmemfs is a memory-based file system, files and directories that are created are not persistent across reboots or unmounts.</p>
EXTENDED DESCRIPTION	<p>To mount the xmemfs file system, do the following: mount -F xmemfs -osize=4g xmem directory</p> <p>You can also mount a xmemfs file system on <code>/xmem</code> at multi-user startup time prior to physical memory becoming fragmented. To do this, add the following line to your <code>/etc/vfstab</code> file: xmem - /xmem xmemfs - yes largebsize,size=4g</p> <p>The xmemfs file system is expressly designed for performance-driven applications (for example, RDBMS) that require large amounts of physical memory. The xmemfs file system provides file system semantics to manage and access extended memory spaces that exceed 4 Gbytes. From an application perspective, extended memory under the control of a mounted xmemfs file system is viewed as a single, large memory pool that can be partitioned as needed through file creation. You can obtain windows into each memory partition by using <code>mmap(2)</code>.</p> <p>Memory controlled by xmemfs can be partitioned by creating files of the required size in the file system. The xmemfs file system allocates sufficient block-sized memory pages for a file based on the file's size. Files can be created using any standard file utility, including <code>mkfile(1M)</code> and <code>dd(1M)</code>. The xmemfs file system optimizes the creation of large files that initially contain all zeroes by allocating memory pages for the file 'hole' that is created by writing beyond the end of file.</p> <p>If sufficient xmemfs extended memory is available, an application can quickly create an 8 Gbyte file in the xmemfs file system by using <code>lseek(2)</code> to offset 8GB-1 and then <code>write(2)</code> a one-byte buffer containing zero. With xmemfs, you can share and protect partitioned memory by setting appropriate file permissions. To avoid wasting memory</p>

resources, (especially with the `-largebsize` option specified), newly created option-specified files should be a multiple of the block size of the `xmemfs` file system. Creation of many small files is strongly discouraged. See `statvfs(2)` for information on determining file system block sizes.

The `xmemfs` file system should only be used with performance-driven applications that require quick access to large amounts of physical memory. Using `xmemfs` for other applications may result in non-optimal use of system resources and possible system performance degradation.

To maximize `xmemfs` ability to access a file's extended memory partition, use `mmap(2)`. The initial `mmap(2)` call enables the system to assign a map size containing as much memory as an application may actively access at any time. The map size is constrained by the application's virtual address space, (usually a maximum of 3 Gbytes on machines with more than 4 Gbytes of physical memory). To access extended memory that is not contained in the existing mapping, use `mmap(2)` with the `-MAP_FIXED` flag to remap a window within the address range returned by the initial `mmap` call.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	i386
Interface Stability	Evolving

SEE ALSO `df(1M)`, `mount(1M)`, `mount_xmemfs(1M)`, `mmap(2)`, `mount(2)`, `statvfs(2)`, `umount(2)`, `vfstab(4)`

DIAGNOSTICS If the `xmemfs` file system runs out of space, the following message is displayed in the console indicating that there is insufficient memory to satisfy a `write(2)` request:

```
directory: File system full, no memory
```

WARNINGS Files and directories on an `xmemfs` file system are not preserved across reboots or unmounts.

zcons(7D)

NAME	zcons – Zone console device driver						
DESCRIPTION	<p>The <code>zcons</code> character driver exports the console for system zones. The driver is comprised of two "sides:" a master side with which applications in the global zone communicate, and a slave side, which receives I/O from the master side. The slave side is available in the global zones.</p> <p>Applications must not depend on the location of <code>/dev</code> or <code>/devices</code> entries exported by <code>zcons</code>. Inside a zone, the <code>zcons</code> slave side is fronted by <code>/dev/console</code> and other console-related symbolic links, which are used by applications that expect to write to the system console.</p> <p>The <code>zcons</code> driver is Sun Private, and may change in future releases.</p>						
FILES	<p><code>/dev/zcons/<zonename>/masterconsole</code> Global zone master side console for zone <code><zonename></code>.</p> <p><code>/dev/zcons/<zonename>/slaveconsole</code> Global zone slave side console for zone <code><zonename></code>.</p> <p><code>/dev/zconsole</code> Non-global zone console (slave side).</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>SUNWcsu</td></tr><tr><td>Interface Stability</td><td>Sun Private</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	SUNWcsu	Interface Stability	Sun Private
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	SUNWcsu						
Interface Stability	Sun Private						
SEE ALSO	<code>zoneadm(1M)</code> , <code>zonecfg(1M)</code> , <code>attributes(5)</code> , <code>zones(5)</code>						

NAME	zero – source of zeroes
DESCRIPTION	<p>A zero special file is a source of zeroed unnamed memory.</p> <p>Reads from a zero special file always return a buffer full of zeroes. The file is of infinite length.</p> <p>Writes to a zero special file are always successful, but the data written is ignored.</p> <p>Mapping a zero special file creates a zero-initialized unnamed memory object of a length equal to the length of the mapping and rounded up to the nearest page size as returned by <code>sysconf</code>. Multiple processes can share such a zero special file object provided a common ancestor mapped the object <code>MAP_SHARED</code>.</p>
FILES	<code>/dev/zero</code>
SEE ALSO	<code>fork(2)</code> , <code>mmap(2)</code> , <code>sysconf(3C)</code>

zs(7D)

NAME	zs – Zilog 8530 SCC serial communications driver
SYNOPSIS	<pre>#include <fcntl.h> #include <sys/termios.h> open("/dev/term/n", mode); open("/dev/tty/n", mode); open("/dev/cua/n", mode);</pre>
DESCRIPTION	<p>The Zilog 8530 provides two serial input/output channels capable of supporting a variety of communication protocols. A typical system uses two or more of these devices to implement essential functions, including RS-423 ports (which also support most RS-232 equipment), and the console keyboard and mouse devices.</p> <p>The <code>zs</code> module is a loadable STREAMS driver that provides basic support for the Zilog 8530 hardware and basic asynchronous communication support. The driver supports the <code>termio(7I)</code> device control functions specified by flags in the <code>c_cflag</code> word of the <code>termios</code> structure and by the <code>IGNBRK</code>, <code>IGNPAR</code>, <code>PARMRK</code>, or <code>INPCK</code> flags in the <code>c_iflag</code> word. All other <code>termio(7I)</code> functions must be performed by STREAMS modules pushed atop the driver. When a device is opened, the <code>ldterm(7M)</code> and <code>ttcompat(7M)</code> STREAMS modules are automatically pushed on top of the stream, providing the standard <code>termio(7I)</code> interface.</p> <p>The character-special devices <code>/dev/term/a</code> and <code>/dev/term/b</code> are used to access the two serial ports on the CPU board.</p> <p>Valid name space entries are <code>/dev/cua/[a-z]</code>, <code>/dev/term/[a-z]</code> and <code>/dev/tty[a-z]</code>. The number of entries used in a name space are machine dependent.</p> <p>The <code>/dev/tty[n]</code> device names only exist if the <i>SunOS 4.x Binary Compatibility Package</i> is installed. The <code>/dev/tty[n]</code> device names are created by the <code>ucblinks</code> command, which is available only with the <i>SunOS 4.x Binary Compatibility Package</i>.</p> <p>To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature is available that is controlled by the minor device number. By accessing character-special devices with names of the form <code>/dev/cua/[n]</code>, it is possible to open a port without the Carrier Detect signal being asserted, either through hardware or an equivalent software mechanism. These devices are commonly known as dial-out lines.</p> <p>Once a <code>/dev/cua/[n]</code> line is opened, the corresponding tty line cannot be opened until the <code>/dev/cua/n</code> line is closed. A blocking open will wait until the <code>/dev/cua/[n]</code> line is closed (which will drop Data Terminal Ready, and Carrier Detect) and carrier is detected again. A non-blocking open will return an error. If the tty line has been opened successfully (usually only when carrier is recognized on the</p>

modem), the corresponding `/dev/cua/[n]` line cannot be opened. This allows a modem to be attached to `/dev/term/[n]` (renamed from `/dev/tty[n]`) and used for dial-in (by enabling the line for login in `/etc/inittab`) and also used for dial-out (by `tip(1)` or `uucp(1C)`) as `/dev/cua/[n]` when no one is logged in on the line.

Note – This module is affected by the setting of specific eeprom variables. For information on parameters that are persistent across reboots, see the `eeprom(1M)` man page.

IOCTLS The `zs` module supports the standard set of `termio ioctl()` calls.

If the `CRTSCTS` flag in the `c_cflag` field is set, output will be generated only if CTS is high; if CTS is low, output will be frozen. If the `CRTSCTS` flag is clear, the state of CTS has no effect.

If the `CRTSXOFF` flag in the `c_cflag` field is set, input will be received only if RTS is high; if RTS is low, input will be frozen. If the `CRTSXOFF` flag is clear, the state of RTS has no effect.

The `termios CRTSCTS` (respectively `CRTSXOFF`) flag and `termiox CTSXON` (respectively `RTSXOFF`) can be used interchangeably.

Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls.

The state of the DCD, CTS, RTS, and DTR interface signals may be queried through the use of the `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` arguments to the `TIOCMGET ioctl` command, respectively. Due to hardware limitations, only the RTS and DTR signals may be set through their respective arguments to the `TIOCMSET`, `TIOCMBIS`, and `TIOCMBIC ioctl` commands.

The input and output line speeds may be set to any of the speeds supported by `termio`. The input and output line speeds cannot be set independently; for example, when you set the the output speed, the input speed is automatically set to the same speed.

When the driver is used to service the serial console port, it supports a `BREAK` condition that allows the system to enter the debugger or the monitor. The `BREAK` condition is generated by hardware and it is usually enabled by default. A `BREAK` condition originating from erroneous electrical signals cannot be distinguished from one deliberately sent by remote DCE. The Alternate Break sequence can be used to remedy this.

Due to a risk of incorrect sequence interpretation, SLIP and certain other binary protocols should not be run over the serial console port when Alternate Break sequence is in effect. Although PPP is a binary protocol, it is able to avoid these sequences using the ACCM feature in *RFC 1662*. For Solaris PPP 4.0, you do this by adding the following line to the `/etc/ppp/options` file (or other configuration files used for the connection; see `pppd(1M)` for details):

```
asynctmap 0x00002000
```

zs(7D)

By default, the Alternate Break sequence is three characters: carriage return, tilde and control-B (CR ~ CTRL-B), but may be changed by the driver. For more information on breaking (entering the debugger or monitor), see `kbd(1)` and `kb(7M)`.

ERRORS

An `open` will fail under the following conditions:

- ENXIO The unit being opened does not exist.
- EBUSY The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open.
- EBUSY The port is in use by another serial protocol.
- EBUSY The unit has been marked as exclusive-use by another process with a `TIOCEXCL ioctl()` call.
- EINTR The open was interrupted by the delivery of a signal.

FILES

- `/dev/cua/[a-z]` dial-out tty lines
- `/dev/term/[a-z]` dial-in tty lines
- `/dev/tty[a-z]` binary compatibility package device names

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

SEE ALSO

`eeeprom(1M)`, `kadb(1M)`, `tip(1)`, `ucblinks(1B)`, `cu(1C)`, `uucp(1C)`, `ports(1M)`, `pppd(1M)`, `ioctl(2)`, `open(2)`, `attributes(5)`, `zsh(7D)`, `termio(7I)`, `kb(7M)`, `ldterm(7M)`, `ttcompat(7M)`

DIAGNOSTICS

- `zsn` : silo overflow.
The Zilog 8530 character input silo overflowed before it could be serviced.
- `zsn` : ring buffer overflow.
The driver's character input ring buffer overflowed before it could be serviced.

NAME	zsh – On-board serial HDLC/SDLC interface
SYNOPSIS	<pre>#include <fcntl.h> open (/dev/zshn, mode); open (/dev/zsh, mode);</pre>
DESCRIPTION	<p>The <code>zsh</code> module is a loadable STREAMS driver that implements the sending and receiving of data packets as HDLC frames over synchronous serial lines. The module is not a standalone driver, but instead depends upon the <code>zs</code> module for the hardware support required by all on-board serial devices. When loaded this module acts as an extension to the <code>zs</code> driver, providing access to an HDLC interface through character-special devices.</p> <p>The <code>zshn</code> devices provide what is known as a data path which supports the transfer of data via <code>read(2)</code> and <code>write(2)</code> system calls, as well as <code>ioctl(2)</code> calls. Data path opens are exclusive in order to protect against injection or diversion of data by another process.</p> <p>The <code>zsh</code> device provides a separate <code>control</code> path for use by programs that need to configure or monitor a connection independent of any exclusive access restrictions imposed by data path opens. Up to three control paths may be active on a particular serial channel at any one time. Control path accesses are restricted to <code>ioctl(2)</code> calls only; no data transfer is possible.</p> <p>When used in synchronous modes, the Z8530 SCC supports several options for clock sourcing and data encoding. Both the transmit and receive clock sources can be set to be the external Transmit Clock (TRxC), external Receive Clock (RTxC), the internal Baud Rate Generator (BRG), or the output of the SCC's Digital Phase-Lock Loop (DPLL).</p> <p>The Baud Rate Generator is a programmable divisor that derives a clock frequency from the PCLK input signal to the SCC. A programmed baud rate is translated into a 16-bit time constant that is stored in the SCC. When using the BRG as a clock source the driver may answer a query of its current speed with a value different from the one specified. This is because baud rates translate into time constants in discrete steps, and reverse translation shows the change. If an exact baud rate is required that cannot be obtained with the BRG, an external clock source must be selected.</p> <p>Use of the DPLL option requires the selection of NRZI data encoding and the setting of a non-zero value for the baud rate, because the DPLL uses the BRG as its reference clock source.</p> <p>A local loopback mode is available, primarily for use by the <code>syncloop(1M)</code> utility for testing purposes, and should not be confused with SDLC loop mode, which is not supported on this interface. Also, an <code>auto-echo</code> feature may be selected that causes all incoming data to be routed to the transmit data line, allowing the port to act as the remote end of a digital loop. Neither of these options should be selected casually, or left in use when not needed.</p>

The zsh driver keeps running totals of various hardware generated events for each channel. These include numbers of packets and characters sent and received, abort conditions detected by the receiver, receive CRC errors, transmit underruns, receive overruns, input errors and output errors, and message block allocation failures. Input errors are logged whenever an incoming message must be discarded, such as when an abort or CRC error is detected, a receive overrun occurs, or when no message block is available to store incoming data. Output errors are logged when the data must be discarded due to underruns, CTS drops during transmission, CTS timeouts, or excessive watchdog timeouts caused by a cable break.

IOCTLS

The zsh driver supports several `ioctl()` commands, including:

<code>S_IOCGETMODE</code>	Return a <code>struct scc_mode</code> containing parameters currently in use. These include the transmit and receive clock sources, boolean loopback and NRZI mode flags and the integer baud rate.
<code>S_IOCSETMODE</code>	The argument is a <code>struct scc_mode</code> from which the SCC channel will be programmed.
<code>S_IOCGETSTATS</code>	Return a <code>struct sl_stats</code> containing the current totals of hardware-generated events. These include numbers of packets and characters sent and received by the driver, aborts and CRC errors detected, transmit underruns, and receive overruns.
<code>S_IOCCLRSTATS</code>	Clear the hardware statistics for this channel.
<code>S_IOCGETSPEED</code>	Returns the currently set baud rate as an integer. This may not reflect the actual data transfer rate if external clocks are used.
<code>S_IOCGETMCTL</code>	Returns the current state of the CTS and DCD incoming modem interface signals as an integer.

The following structures are used with zsh `ioctl()` commands:

```

struct scc_mode {
    char sm_txclock; /* transmit clock sources */
    char sm_rxclock; /* receive clock sources */
    char sm_iflags; /* data and clock inversion flags (non-zsh) */
    uchar_t sm_config; /* boolean configuration options */
    int sm_baudrate; /* real baud rate */
    int sm_retval; /* reason codes for ioctl failures */
};
struct sl_stats {
    long ipack; /* input packets */
    long opack; /* output packets */
    long ichar; /* input bytes */
    long ochar; /* output bytes */
    long abort; /* abort received */
    long crc; /* CRC error */
    long cts; /* CTS timeouts */
    long dcd; /* Carrier drops */
}

```



```

long overrun;      /* receive overrun */
long underrun;     /* transmit underrun */
long ierror;       /* input error */
long oerror;       /* output error */
long nobuffers;    /* receive side memory allocation failure */
};

```

ERRORS An `open()` will fail if a STREAMS message block cannot be allocated, or:

ENXIO The unit being opened does not exist.

EBUSY The device is in use by another serial protocol.

An `ioctl()` will fail if:

EINVAL An attempt was made to select an invalid clocking source.

EINVAL The baud rate specified for use with the baud rate generator would translate to a null time constant in the SCC's registers.

FILES `/dev/zsh[0-1]`, `/dev/zsh` character-special devices
`/usr/include/sys/ser_sync.h` header file specifying synchronous serial communication definitions

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86

SEE ALSO `syncinit(1M)`, `syncloop(1M)`, `syncstat(1M)`, `ioctl(2)`, `open(2)`, `read(2)`, `write(2)`, `attributes(5)`, `zs(7D)`

Refer to the *Zilog Z8530 SCC Serial Communications Controller Technical Manual* for details of the SCC's operation and capabilities.

DIAGNOSTICS `zsh data open failed, no memory, rq=nnn`

`zsh clone open failed, no memory, rq=nnn`

A kernel memory allocation failed for one of the private data structures. The value of `nnn` is the address of the read queue passed to `open(2)`.

`zsh_open: can't alloc message block`

The open could not proceed because an initial STREAMS message block could not be made available for incoming data.

`zsh: clone device d must be attached before use!`

An operation was attempted through a control path before that path had been attached to a particular serial channel.

`zshn: invalid operation for clone dev.`

An inappropriate STREAMS message type was passed through a control path. Only `M_IOCTL` and `M_PROTO` message types are permitted.

zsh(7D)

zshn: not initialized, can't send message

An `M_DATA` message was passed to the driver for a channel that had not been programmed at least once since the driver was loaded. The SCC's registers were in an unknown state. The `S_IOCSETMODE` ioctl command performs the programming operation.

zshn: transmit hung

The transmitter was not successfully restarted after the watchdog timer expired.

NAME zulu – Sun XVR-4000 Graphics Accelerator driver

DESCRIPTION The zulu driver is the device driver for the Sun XVR-4000 Graphics Accelerator.

FILES /dev/fbs/zulun
device special file

/usr/lib/zulu.unicode
zulu microcode

ATTRIBUTES See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWzulux

SEE ALSO SUNWzulu_config(1M), zuludaemon(1M), attributes(5)

zulu(7d)

Index

Numbers and Symbols

- 1394 OpenHCI host controller driver — hci1394, 229
- 24-bit UPA color frame buffer and graphics accelerator — ffb, 201
- 3COM EtherLink III PCMCIA Ethernet Adapter — pcelx, 436
- 3Com Ethernet device driver — elxl, 162
- 4BSD compatibility module — ttcompat, 681

A

- aac— SCSI HBA driver for Adaptec AdvancedRAID Controller, 22
- Acer Laboratories Inc. M5451 audio processor interface — audiots, 63
- Adaptec Ultra320 SCSI host bus adapter driver — adpu320, 25
- Address Resolution Protocol, *See* ARP
- adp — low-level module for controllers based on Adaptec AIC-7870P and AIC-7880P SCSI chips, 23
- adpu320 — Adaptec Ultra320 SCSI host bus adapter driver, 25
- afb— Elite3D graphics accelerator driver, 27
- ifb— IFB graphics accelerator driver, 266
- Alarm device driver — tsalarm, 677
- tsalarm — Alarm device driver, 677
- Amd 8111 south bridge audio digital controller — audio810, 51
- AMD PCnet Ethernet controller device driver — pcn, 447
- amr— SCSI HBA driver for LSI MegaRAID Controller, 28
- an I2O OS sepcific module that supports SCSSA interface, an I2O, 248
- arp — Address Resolution Protocol, 29
- arp ioctl
 - SIOCDDARP — delete arp entry, 29
 - SIOCGARP — get arp entry, 29
 - SIOCSARP — set arp entry, 29
- asy — asynchronous serial port driver, 33
- asynchronous serial port driver — asy, 33
- asynchronous serial port driver — su, 620
- AT attachment disk driver — ata, 36
- driver — cpuid, 109
- ata — AT attachment disk driver, 36
- audio — audio device interface, 40
- audio device, Sound Blaster 16/Pro/AWE32 — sbpro, 523
- audio_support — audio device independent driver interface, 61
- audio810 — Amd 8111 south bridge audio digital controller interface, 51
- audio810 —Amd 8111 south bridge audio digital controller, 51
- audio810 — Amd 8111 south bridge audio digital controller interface
 - Driver Version, 51
 - Sample Granularity, 51
- audiocs — Crystal Semiconductor 4231 Audio driver, 54
- audiocs — Crystal Semiconductor 4231 audio Interface

- audiocs — Crystal Semiconductor 4231 audio Interface (Continued)
 - Audio Data Formats for the Multimedia 4231 Codec, 54
- audio810 — Amd 8111 south bridge audio digital controller, Audio Interfaces, 51
- audiocs — Crystal Semiconductor 4231 Audio driver, Audio Interfaces, 54
- audiots — Acer Laboratories Inc. M5451 audio processor interface
 - Audio Interfaces, 63
- audiocs — Crystal Semiconductor 4231 audio Interface
 - Audio Status Change Notification, 55
 - Driver Version, 54
 - Sample Granularity, 54
 - Setting Audio Mixer Mode, 54
- audiots — Acer Laboratories Inc. M5451 audio processor interface, 63
 - Audio Status Change Notification, 64
 - Driver Version, 63
 - Sample Granularity, 63
- av1394— 1394 audio/video driver, 66
- 1394 audio/video driver — av1394, 66

B

- bbc_beep — Platform-dependent Beep driver for BBC-based hardware, 67
- bd — SunButtons and SunDials STREAMS module, 68
- bge — Broadcom BCM57xx Gigabit Ethernet driver, 70
- bpp — bi-directional parallel port, 74
- Broadcom BCM57xx Gigabit Ethernet driver — bge, 70
- bufmod — STREAMS Buffer module, 80
- built-in mouse device interface — kdmouse, 346

C

- cadp — SCSI host bus adapter driver, 85
- cadp160 — SCSI host bus adapter driver, 84
- CD-ROM — ISO 9660 CD-ROM filesystem — hfsf, 241

- cdio— CD-ROM control operations, 89
- CDROM control operations —cdio, 89
- ce — ce Gigabit-Ethernet device driver, 97
 - ce Primitives, 99
- ce Gigabit-Ethernet device driver — ce, 97
- cgsix — accelerated 8-bit color frame buffer, 101
- chs — IBM ServeRAID PCI host adapter driver, 102
- Data Layer Class of Service Marker— dlcosmk, 136
- cmdk — common disk driver, 103
- Differentiated Services Code Point Marker— dscpmk, 147
- Cogent EM960/EM100 Ethernet controller device driver — dnet, 142
- color graphics interface, accelerated 8-bit color frame buffer — cgsix, 101
- common disk driver — cmdk, 103
- Single and Two Rate Three Conformance Level Meter— tokenmt, 673
- connections, unique stream, line discipline — connld, 104
- connld — line discipline for unique connections, 104
- console devices, sysmsg, 629
- console, STREAMS-based console interface, 106
- contract file system — ctfs, 111
- converts mouse protocol to Firm Events —
 - vuidmice, 758
 - vuid2ps2, 758
 - vuid3ps2, 758
 - vuidm3p, 758
 - vuidm4p, 758
 - vuidm5p, 758
- core memory, image — mem, 381
- cpr — suspend and resume module, 107
- cpuid — driver, 109
- Crystal Semiconductor 4231 Audio driver — audiocs, 54
- ctfs — contract file system, 111
- ctsmc — System Management Controller driver, 112
- cvc — virtual console driver, 113
- cvcredir — virtual console redirection driver, 114

D

- D-LINK Ethernet controller device driver — dnet, 142
- dad — driver for IDE disk devices, 115
 - Device Statistics Support, 115
- dapl — Tavor uDAPL service driver, 118
- Data Link Provider Interface, — dlpi, 137
- Davicom Fast Ethernet driver for Davicom DM9102A — dmfe, 140
- dbri — ISDN and audio interface, 119
- dbri — Dual Basic Rate ISDN and audio Interface
 - Audio Data Formats for BRI Interfeces, 120
 - Audio Data Formats for the Multimedia Codec/SpeakerBox, 119
 - Audio Interfaces, 119
 - Audio Ports, 121
 - Audio Status Change Notification, 122
 - ISDN Interfaces, 119
 - Sample Granularity, 122
- delete arp entry ioctl — SIOCDARP, 29
- devfs— devices file system, 124
- device interface, Microsoft Bus Mouse — msm, 403
- devices file system, devices file system — devfs, 124
- usbser_edge — Digi Edgeport USB to serial converter driver, 740
- disk control operations — dkio, 126
- disk driver, fd — floppy, 192
- disk quotas — quotactl(), 498
- dkio — disk control operations, 126
- dlcosmk —Data Layer Class of Service Marker, 136
- dlpi — Data Link Provider Interface, 137
- dmfe — Davicom Fast Ethernet driver for Davicom DM9102A, 140
- dnet — DEC 21040/21140-based Ethernet Controllers, 142
- DOS, DOS formatted file system — pcfs, 438
- dpt — DPT ServeRAID IV SCSI host bus adapter and RAID adapter driver, 144
- DPT ServeRAID IV SCSI host bus adapter and RAID adapter driver — dpt, 144
- dynamic reconfiguration driver — Sun Enterprise 10000, 146
- Sun Enterprise 10000 dynamic reconfiguration driver, 146

- driver for fast poll on many file descriptors — poll, 474
- driver for IDE disk devices — dad, 115
- driver for parallel port — lp, 371
- drivers for floppy disks and floppy disk controllers — fd, 192
 - fdc, 192
- drivers, SCSI tape devices — st, 588
- dscpmk —Differentiated Services Code Point Marker, 147
- dtrace— DTrace dynamic tracing facility, 148
- DTrace dynamic tracing facility — dtrace, 148
- DTrace kernel lock instrumentation provider — lockstat, 363
- DTrace profile interrupt provider — profile, 483
- DTrace statically defined tracing provider — sdt, 549
- DTrace function boundary tracing provider — fbt, 186
- DTrace user instruction tracing provider — fasttrap, 183
- Dual Basic Rate ISDN and audio Interface — dbri, 119

E

- e1000g — Gigabit Ethernet device driver, 149
- ecpp — bi-directional parallel port, 154
- ehci— Enhanced host controller driver, 160
- Elite3D graphics accelerator driver — afb, 27
- elxl — 3Com Ethernet device driver, 162
- Enhanced host controller driver — ehci, 160
- eri — eri Fast-Ethernet device driver, 165
- esp — ESP SCSI Host Bus Adapter Driver, 169
- ESP SCSI Host Bus Adapter Driver — esp, 169
- Ethernet device driver, SMC EtherPower II 10/100 (9432) Ethernet device driver — spwr, 583
- extended memory file system — xmemfs, 776

F

- Fair share scheduler — FS, 205
- FS — Fair share scheduler, 205
- fas — FAS SCSI Host Bus Adapter Driver, 175

FAS SCSI Host Bus Adapter Driver — *fas*, 175
 fasttrap— DTrace user instruction tracing provider, 183
 fbio — frame buffer control operations, 184
 fbt— DTrace function boundary tracing provider, 186
 fcip — IP/ARP over Fibre Channel datagram encapsulation driver, 187
 fcp — Fibre Channel protocol driver, 190
 fctl — Sun Fibre Channel transport library, 191
 fd — drivers for floppy disks and floppy disk controllers, 192
 fdc — drivers for floppy disks and floppy disk controllers, 192
 FDGETCHANGE — get status of disk changed, 197
 fdio — disk control operations, 197
 FDIOGCHAR — get floppy characteristics, 197
 FDIOGCHAR — set floppy characteristics, 197
 FDKEJECT — eject floppy, 197
 ffb — 24-bit UPA color frame buffer and graphics accelerator, 201
 Sun Fibre Channel port driver — *fp*, 203
 Fibre Channel protocol driver — *fcp*, 190
 file system, *quotactl()* — disk quotas, 498
 floppy disk driver — *fd*, 192
 floppy disk control operations — *fdio*, 197
 Flow Accounting module— *flowacct*, 202
flowacct —Flow Accounting module, 202
fp — Sun Fibre Channel port driver, 203
 frame buffer control operations, — *fbio*, 184

G

ge — Gigabit-Ethernet device driver, 208
ge Primitives, 210
 Gigabit-Ethernet device driver — *ge*, 208
 general properties of Internet Protocol network interfaces — *if*, 271
 general properties of Internet Protocol network interfaces — *if_tcp*, 271
 General purpose I/O driver for SuperIO — *gpio_87317*, 227
 general terminal interface — *termio*, 638
 Generic LAN Driver — *gld*, 212
 Generic printer interface — *prnio*, 479
prnio — Generic printer interface, 479

Generic SCSI device driver — *sgen*, 562
gld — Generic LAN Driver, 212
 FDDI SNAP processing, 213
 Implemented DLPI Primitives, 214
 Implemented *ioctl* Functions, 216
 Network Statistics, 218
 Style 1 and 2 Providers, 214
 Token Ring SNAP processing, 213
 Token Ring Source Routing, 214
glm — GLM SCSI Host Bus Adapter Driver, 222
 GLM SCSI Host Bus Adapter Driver — *glm*, 222
gpio_87317— General purpose I/O driver for SuperIO, 227
grbeep — Platform-dependent Beep driver for SMBus-based hardware, 228

H

hci1394 — 1394 OpenHCI host controller driver, 229
hdio — SMD and IPI disk control operations, 230
 — Human Interface Device (HID) class driver, 232
 High Sierra filesystem, *See* *hsfs*
hme — SUNW,hme Fast-Ethernet device driver, 234
hme Primitives, 236
hsfs, filesystem — *hsfs*, 241
hubd — USB hub driver, 244
 Human Interface Device (HID) class driver — *hid*, 232

I

I/O
 data link provider interface — *dlpi*, 137
 extended terminal interface — *termiox*, 659
ioctls that operate directly on sockets — *sockio*, 581
 STREAMS *ioctl* commands — *streamio*, 604
ib — InfiniBand Bus Nexus Driver, 249
ibcm — Solaris InfiniBand Communication Manager, 252

ibd — Infiniband IPoIB device driver, 253
 ibdm — Solaris InfiniBand Device Manager, 256
 IBM ServeRAID PCI host adapter driver — chs, 102
 ibmf — InfiniBand Management Transport Framework, 257
 ibtl — Solaris InfiniBand Transport Layer, 258
 iButton smart card terminal driver — ocf_ibutton, 426
 icmp — Internet Control Message Protocol, 261
 icmp6 — Internet Control Message Protocol for Internet Protocol Version 6, 259
 IDE Host Bus Adapter Driver — uata, 695
 if — general properties of Internet Protocol network interfaces, 271
 if_tcp — general properties of Internet Protocol network interfaces, 271
 Application Programming Interface, 271
 IFB graphics accelerator driver — ifb, 266
 inet — Internet protocol family, 285
 inet6 — Internet protocol family for Internet Protocol version 6, 280
 InfiniBand Bus Nexus Driver — ib, 249
 Solaris InfiniBand Communication Manager — ibcm, 252
 Solaris InfiniBand Device Manager — ibdm, 256
 Infiniband IPoIB device driver — ibd, 253
 InfiniBand Management Transport Framework — ibmf, 257
 InfiniHost MT23108 InfiniBand (IB) Driver — tavor, 631
 Solaris InfiniBand Transport Layer — ibtl, 258
 Intel D100 Ethernet device driver — iprb, 305
 Intel Gigabit Ethernet device driver — e1000g, 149
 Intel i82365SL PC Card Interface Controller — pcic, 443
 Intel Ethernet device driver, Intel — e1000g, 149
 Intel Ethernet device driver, Intel — iprb, 305
 Internet Control Message Protocol — icmp, 261
 Internet Protocol — ip, 295
 Internet Transmission Control Protocol — tcp, 633
 Internet User Datagram Protocol — udp, 699
 Internet Control Message Protocol for Internet Protocol Version 6 — icmp6, 259
 Internet Protocol, to Ethernet addresses — arp, 29
 Internet protocol family — inet, 285
 Internet protocol family for Internet Protocol version 6 — inet6, 280
 Internet Protocol Security Architecture — ipsec, 308
 Internet Protocol Version 6 — ip6, 288
 Internet Protocol Version 6 — ip6, 288
 ioctl's for sockets
 SIOCDDARP — delete arp entry, 29
 ioctl's for sockets, SIOCGARP — get arp entry, 29
 ioctl's for Internet socket descriptors, SIOCSARP — set arp entry, 29
 ioctl's for floppy
 FDEJECT — eject floppy, 197
 FDGETCHAGE — get status of disk changed, 197
 FDIOCHAR — get floppy characteristics, 197
 FDIOCHAR — set floppy characteristics, 197
 ioctl's for sockets
 SIOCADDRT — add route, 516
 SIOCDELRT — delete route, 516
 ioctl's for terminals
 TIOCPKT — set/clear packet mode (pty), 490
 TIOCREMOTE — remote input editing, 490
 TIOCSTART — start output (like control-Q), 490
 TIOCSTOP — stop output (like control-S), 490
 ip — Internet Protocol, 295
 IP/ARP over Fibre Channel datagram encapsulation driver — fcip, 187
 IP Generic Packet Classifier — ipgpc, 301
 IP quality of service — ipqos, 303
 ip6 — Internet Protocol Version 6, 288
 ipgpc — IP Generic Packet Classifier, 301
 ipqos — IP quality of service, 303
 iprb — D100 Ethernet device driver, 305
 ipsec — Internet Protocol Security Architecture, 308
 isdnio — generic ISDN interface, 314

ISO 9660 — ISO 9660 CD-ROM filesystem —
hsfs, 241
isp — ISP SCSI Host Bus Adapter Driver, 327
ISP SCSI Host Bus Adapter Driver — isp, 327
ISP2200/ ISP2300 Family Fibre Channel host
bus adapter driver — qlc, 496

J

jfca — JNI Fibre Channel Adapter (FCA)
Driver, 336
JNI Fibre Channel Adapter (FCA) Driver —
jfca, 336

K

kb — keyboard STREAMS module, 337
Keyboard Compatibility Mode, 341
Keyboard Translation Mode, 337
Keyboard Translation-Table Entries, 337
kdmouse — built-in mouse device
interface, 346
kernel debugger — kmdb, 347
Kernel object filesystem — objfs, 425
kernel packet forwarding database —
route, 512
kernel statistics driver — kstat, 348
kernel symbols — ksyms, 349
keyboard STREAMS module — kb, 337
keyboard STREAMS module for Sun USB
Keyboard — usbkbm, 729
kmdb — kernel debugger, 347
kstat — kernel statistics driver, 348
kysms — kernel symbols, 349

L

LAN support module — gld, 212
ldterm — line discipline for STREAMS terminal
module, 351
line discipline for unique stream connections
— connlld, 104
llc1 — Logical Link Control Protocol Class 1
Driver, 354

lockstat — DTrace kernel lock instrumentation
provider, 363
lofi — loopback file driver, 364
lofs — loopback virtual file system, 365
log — interface to STREAMS error logging and
event tracing, 366
logi — LOGITECH bus mouse device
interface, 370
Logical Link Control Protocol Class 1 Driver —
llc1, 354
LOGITECH Bus Mouse device interface —
logi, 370
loopback file driver — lofi, 364
loopback transport providers
— ticlts, 665
— ticots, 665
— ticotsord, 665
loopback virtual file system — lofs, 365
low-level module for controllers based on
Adaptec AIC-7870P and AIC-7880P SCSI
chips — adp, 23
low-level module for the AMD PCscsi, PCscsi II,
PCnet-SCSI, and Qlogic QLA510 PCI-to-SCSI
bus adapters — pcscsi, 452
lp — driver for parallel port, 371

M

m64 — frame buffers device driver, 373
magnetic tape interface, — mtio, 405
md — pseudo device driver for
metadevice, 374
mem— image of core memory, 381
memory, core, image — mem, 381
memory, zeroed unnamed, source — zero, 779
memory based filesystem — tmpfs, 671
message output collection, msglog, 402
mhd — multihost disk control operations, 383
Automatic Probing, 386
MHIOCENFAILFAST Ioctl, 386
Non-shared multihost disks:, 383
Shared Multihost Disks, 384
Microsoft Bus Mouse device interface —
msm, 403
mixer — audio mixer device interface, 388
monitor, PROM monitor configuration interface
— openprom, 429

Mouse device interface, LOGITECH Bus Mouse device interface — logi, 370
 mpt — SCSI host bus adapter driver, 398
 Driver Configuration, 398
 msglog, message output collection, 402
 msm — Microsoft Bus Mouse device interface, 403
 mtio — general magnetic tape interface, 405
 multihost disk control operations — mhd, 383

N

llc2— llc2 Class II logical link control driver, 357
 ncrs — SCSI host bus adapter driver, 417
 ncrs — ncrs SCSI Host Bus Adapter Driver, Driver Configuration, 418
 network packet routing device — routing, 516
 null — null file, 424

O

objfs — Kernel object filesystem, 425
 ocf_ibutton — iButton smart card terminal driver, 426
 ohci — OpenHCI host controller driver, 427
 OpenHCI host controller driver — ohci, 427
 openprom — PROM monitor configuration interface, 429

P

packet routing device — routing, 516
 packet routing ioctl
 SIOCADDRT — add route, 516
 SIOCDELRT — delete route, 516
 parallel port, bi-directional — bpp, 74
 parallel port, bi-directional — ecpp, 154
 parallel port, driver for parallel port — lp, 371
 pcata — PCMCIA ATA card device driver, 434
 pcelx — 3COM EtherLink III PCMCIA Ethernet Adapter, 436
 pcfs — DOS formatted file system, 438
 pcic — Intel i82365SL PC Card Interface Controller, 443

pckt — STREAMS Packet Mode module, 445
 PCMCIA ATA card device driver — pcata, 434
 PCMCIA memory card nexus driver — pcmem, 446
 PCMCIA RAM memory card device driver — pcram, 449
 PCMCIA serial card device driver — pcser, 453
 pcmem — PCMCIA memory card nexus driver, 446
 pcn — AMD PCnet Ethernet controller device driver, 447
 pcram — PCMCIA RAM memory card device driver, 449
 pcscsi — low-level module for the AMD PCscsi, PCscsi II, PCnet-SCSI, and Qlogic QLA510 PCI-to-SCSI bus adapters, 452
 pcser — PCMCIA serial card device driver, 453
 pf_key — security association database, 456
 pfb — Sun XVR-100 Graphics Accelerator device driver, 455
 pfmod — STREAMS packet filter module, 466
 pipemod — STREAMS pipe flushing module, 469
 Platform-dependent Beep driver for BBC-based hardware — bbc_beep, 67
 Platform-dependent Beep driver for SMBus-based hardware — grbeep, 228
 Platform Specific Module (PSM) for Tricord Systems Enterprise Server Models ES3000, ES4000 and ES5000 — tpf, 676
 pm — Power Management Driver, 470
 poll — driver for fast poll on many file descriptors, 474
 Power Management Driver — pm, 470
 PPP tunneling pseudo-driver — sppptun, 582
 profile — DTrace profile interrupt provider, 483
 PROM, monitor configuration interface — openprom, 429
 pseudo-terminal driver — pty, 489
 Pseudo Terminal Emulation module, STREAMS — ptem, 484
 ptem — STREAMS Pseudo Terminal Emulation module, 484
 ptm — STREAMS Buffer module, 485
 pts — STREAMS pseudo-tty slave driver, 487
 pty — pseudo-terminal driver, 489

Q

ixgb — SUNWixgb, 10G Gigabit Ethernet driver for Intel 82597ex and Sun Kirkwood Adapter, 333
qfe — SUNW,qfe Quad Fast-Ethernet device driver, 492
qfe Primitives, 494
qlc — ISP2200/ ISP2300 Family Fibre Channel host bus adapter driver, 496
quotactl() — disk quotas, 498
QUS — SCSI ISP10160 Host Bus Adapter Driver, 500

R

ramdisk — ramdisk device driver, 507
ramdisk device driver — ramdisk, 507
random — strong random number generator device, 509
rarp — Reverse address resolution protocol, 511
Realtek 8139 Ethernet device driver — rtls, 518
Realtek Ethernet device driver, Realtek — rtls, 518
remote input editing ioctl — TIOCREMOTE, 490
WCI Remote Shared Memory (WRSM) DLPI driver — wrsmd, 764
Rockwell 2200 SNAP Streams Driver — sxp, 623
route — kernel packet forwarding database, 512
Messages, 513
routing — local network packet routing, 516
routing ioctls
SIOCADDRT — add route, 516
SIOCDELRT — delete route, 516

S

sbpro — Creative Labs Sound Blaster audio device, 523
scsa1394 — SCSI to 1394 bridge driver, 529
scsa2usb — SCSI to USB bridge host bus adapter driver, 532

SCSI enclosure services device driver — ses, 557
SCSI HBA driver for Adaptec AdvancedRAID Controller — aac, 22
SCSI HBA driver for LSI MegaRAID Controller — amr, 28
Class II logical link control driver — llc2, 357
ncrs SCSI host bus adapter driver — ncrs, 417
SCSI host bus adapter driver — cadp160, 84
SCSI host bus adapter driver — cadp, 85
SCSI host bus adapter driver — mpt, 398
SCSI tape devices, driver — st, 588
SCSI to 1394 bridge driver driver — scsa1394, 529
SCSI to USB bridge host bus adapter driver — scsa2usb, 532
SCSI virtual host controller interconnect driver — scsi_vhci, 537
sctp — Stream Control Transmission Protocol, 539
SCSI disk and ATAPI/SCSI CD-ROM device driver — sd, 543
sd — SCSI disk and ATAPI/SCSI CD-ROM device driver, 543
sdt — DTrace statically defined tracing provider, 549
se — Siemens 82532 ESCC serial communications driver, 550
security association database — pf_key, 456
serial communications driver — zs, 780
Serial Optical Controller for Fibre Channel Arbitrated Loop (SOC+) device driver — social, 579
Service Location Protocol — slp, 576
ses — SCSI enclosure services device driver, 557
set/clear, packet mode (pty) ioctl — TIOCPKT, 490
sgen — Generic SCSI device driver, 562
zsh — On-board serial HDLC interface, 554, 783
Shared Memory (WRSM) device driver — wrsm, 762
Siemens 82532 ESCC serial communications driver — se, 550
SIOCDDARP — delete arp entry, 29
SIOCGARP — get arp entry, 29
SIOCSARP — set arp entry, 29

- sip — SIP proxy/registrat/redirect server, 567
- SIP Proxy/registrat/redirect server — sip, 567
- slp — Service Location Protocol, 576
- smbus — System Management Bus controller driver, 578
- SMC Ethernet device drivers, spwr — SMC EtherPower II 10/100 (9432) Ethernet device driver, 583
- SMC EtherPower 8432BT Ethernet controller device driver — dnet, 142
- SMC EtherPower II 10/100 (9432) Ethernet device driver — spwr, 583
- SMD and IPI disk control operations — hdio, 230
- socal — Serial Optical Controller for Fibre Channel Arbitrated Loop (SOC+) device driver, 579
- sockio — ioctls that operate directly on sockets, 581
- sockets, ioctls that operate directly — sockio, 581
- Solaris USB architecture — usba, 721
- Solaris VISUAL I/O control operations, 749
- Sound Blaster 16/Pro/AWE32 audio devices — sbpro, 523
- spwr — SMC EtherPower II 10/100 (9432) Ethernet device driver, 583
- ssd — Fibre Channel Arbitrated Loop disk device driver, 584
- st — driver for SCSI tape devices, 588
- start output (like control-Q) ioctl — TIOCSTART, 490
- stop output (like control-S) ioctl — TIOCSTOP, 490
- STP 4020 PCMCIA Adapter, STP 4020 PCMCIA Adapter, 603
- stp4020 — STP 4020 PCMCIA Adapter, 603
- Stream Control Transmission Protocol — sctp, 539
- STREAM virtual network interface driver — grbeep, 755
- STREAMS
 - interface to error logging — log, 366
 - interface to event tracing — log, 366
 - line discipline for unique stream connections — connld, 104
 - loopback transport providers — ticlts, ticots, ticotsord, 665
- STREAMS (Continued)
 - On-board serial HDLC interface — se_hdlc, 554
 - On-board serial HDLC interface — zsh, 783
 - standard terminal line discipline module — ldterm, 351
 - Transport Interface cooperating module — timod, 667
 - Transport Interface read/write interface module — tirdwr, 669
 - V7, 4BSD, XENIX compatibility module — ttcompat, 681
- STREAMS Administrative Driver — sad, 520
- STREAMS-based console interface, console, 106
- STREAMS-based serial I/O multiplexer interface, ttymux, 690
- STREAMS Buffer module — bufmod, 80
- STREAMS Buffer module — ptm, 485
- STREAMS ioctl commands — streamio, 604
- STREAMS module, SunButtons and SunDials — bd, 68
- STREAMS Packet Filter Module — pfmod, 466
- STREAMS Packet Mode module — pckt, 445
- STREAMS pipe flushing module — pipemod, 469
- STREAMS Pseudo Terminal Emulation module — ptem, 484
- STREAMS pseudo-tty slave driver — pts, 487
- strong random number generator device — random, 509
- su — asynchronous serial port driver, 620
- Sun Fibre Channel transport library — fctl, 191
- Sun XVR-4000 Graphics Accelerator driver — zulu, 787
- SunButtons and SunDials STREAMS module — bd, 68
- SUNW,hme Fast-Ethernet device driver — hme, 234
- SUNW,qfe Quad Fast-Ethernet device driver — qfe, 492
- eri Fast-Ethernet device driver — eri, 165
- SUNWixgb, 10G Gigabit Ethernet driver for Intel 82597ex and Sun Kirkwood Adapter — ixgb, 333
- suspend and resume module — cpr, 107
- xsp — Rockwell 2200 SNAP Streams Driver, 623
- symhisl — SCSI Host Bus Adapter Driver, 626

SCSI Host Bus Adapter Driver — symhisl, 626
sysmsg
 console devices, 629
 system message routing, 629
System Management Bus controller driver —
 smbus, 578
System Management Controller driver—
 ctsmc, 112
system message routing, sysmsg, 629

T

tape, magnetic interface, — mtio, 405
tape interface — mt, 404
tavor — InfiniHost MT23108 InfiniBand (IB)
 Driver, 631
tcp — Internet Transmission Control
 Protocol, 633
terminal, standard STREAMS, line discipline
 module — ldterm, 351
terminal interface, — termio, 638
terminal interface, extended, — termiox, 659
terminal parameters — termiox, 659
termio — general terminal interface, 638
 Canonical mode input processing, 639
 Comparison of the different cases of MIN,
 TIME interaction, 641
 Control Modes, 648
 Default values, 656
 Input modes, 645
 Local modes, 652
 Minimum and Timeout, 654
 Modem disconnect, 643
 Modem lines, 655
 Non-canonical mode input processing, 640
 Output modes, 646
 Special Characters, 642
 Terminal parameters, 644
 Terminal size, 654
 Termio structure, 654
 Writing characters, 642
termiox — extended general terminal
 interface, 659
Time Sliding Window Three Conformance Level
 Meter— tswtclmt, 680
ticlts — loopback transport provider, 665
ticots — loopback transport provider, 665

ticotsord — loopback transport provider, 665
timod — Transport Interface cooperating
 module, 667
timod — Transport Interface read/write
 interface module, 669
TIOCPKT — set/clear packet mode (pty), 490
TIOCREMOTE — remote input editing, 490
TIOCSTART — start output (like
 control-Q), 490
TIOCSTOP — stop output (like control-S), 490
tmpfs — memory based filesystem, 671
tokenmt — Single and Two Rate Three
 Conformance Level Meter, 673
tpf — Platform Specific Module (PSM) for
 Tricord Systems Enterprise Server Models
 ES3000, ES4000 and ES5000, 676
Transport Interface cooperating STREAMS
 module — timod, 667
Transport Interface read/write interface
 STREAMS module — timod, 669
tswtclmt — Time Sliding Window Three
 Conformance Level Meter, 680
ttcompat — V7, 4BSD and XENIX STREAMS
 compatibility module, 681
tty — controlling terminal interface, 689
ttymux, STREAMS-based serial I/O multiplexer
 interface, 690

U

uata — IDE Host Bus Adapter Driver, 695
Tavor uDAPL service driver — daplt, 118
udfs— universal disk format file system, 698
udp — Internet User Datagram Protocol, 699
ufs, filesystem — ufs, 701
ugen — USB generic driver, 704
universal disk, universal disk format file system
 — udfs, 698
universal serial optical controller for Fibre
 Channel arbitrated loop (SOC+) device
 driver — usoc, 747
unnamed zeroed memory, source — zero, 779
audioens — Ensoniq 1371/1373 and Creative
 Labs 5880 driver support, 57
usb_ac — USB audio control driver, 723
usb_as — USB audio control driver, 727

USB audio HID STREAMS module —
usb_ah, 726
USB generic driver — ugen, 704
USB hub driver — hubd, 244
usb_mid — USB Multi Interface Driver, 731
USB mouse STREAMS module — usbms, 732
USB Multi Interface Driver — usb_mid, 731
usbprn — USB printer class driver, 735
usba — Solaris USB architecture, 721
usbkbm — keyboard STREAMS module for Sun
USB Keyboard, 729
usbkbm — keyboard STREAMS module for
Sun, IOCTLS, 729
usbms — USB mouse STREAMS module, 732
uscsi — user SCSI command interface, 743
user SCSI command interface — uscsi, 743
usoc — universal serial optical controller for
Fibre Channel arbitrated loop (SOC+) device
driver, 747

V

V7 compatibility module — ttcompat, 681
virtual console driver — cvc, 113
virtual console redirection driver —
cvcredir, 114
vni — STREAM virtual network interface
driver, 755
volfs — Volume Management file system, 756
Volume Management, file system — volfs, 756
vuid2ps2 — converts mouse protocol to Firm
Events, 758
vuid3ps2 — converts mouse protocol to Firm
Events, 758
vuidm3p — converts mouse protocol to Firm
Events, 758
vuidm4p — converts mouse protocol to Firm
Events, 758
vuidm5p — converts mouse protocol to Firm
Events, 758
vuidmice — converts mouse protocol to Firm
Events, 758

W

workstation console — wscons, 765

wrsm — Shared Memory (WRSM) device
driver, 762
wrsmd — WCI Remote Shared Memory
(WRSM) DLPI driver, 764

X

XENIX compatibility module — ttcompat, 681
xmemfs — extended memory file system, 776
Sun XVR-100 Graphics Accelerator device
driver— pfb, 455
jfb — XVR-1200 Graphics Accelerator device
driver, 335
XVR-1200 Graphics Accelerator device driver,
jfb, 335

Z

zcons — Zone console device driver, 778
zero — source of zeroes, 779
Zilog 8530 SCC serial communications driver —
zs, 780
Zone console device driver — zcons, 778
zs — zilog 8530 SCC serial communications
driver, 780
zulu — Sun XVR-4000 Graphics Accelerator
driver, 787

