

VERITAS Volume Manager 4.1

Hardware Notes

Solaris

N13111F

March 2005

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 2005 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo and VERITAS Storage Foundation are trademarks or registered trademarks of VERITAS Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000 Fax 650-527-2908
www.veritas.com

Third-Party Legal Notices

Data Encryption Standard (DES) Copyright

Copyright © 1990 Dennis Ferguson. All rights reserved.

Commercial use is permitted only if products that are derived from or include this software are made available for purchase and/or use in Canada. Otherwise, redistribution and use in source and binary forms are permitted.

Copyright 1985, 1986, 1987, 1988, 1990 by the Massachusetts Institute of Technology. All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided as is without express or implied warranty.

Sun Microsystems Trademarks

Sun, Solaris, SunOS, Java, Sun Java System Cluster, Sun StorEdge, Solstice DiskSuite, Sun Fire, Sun Enterprise, Online: Backup, and Netra are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Contents

Preface	vii
Conventions	viii
Getting Help	ix
Documentation Feedback	ix
Hardware Support in VxVM 4.1	1
Disk Arrays Supported for DMP	2
Vendor ASL packages for Supported Arrays	3
Arrays Supported in a Fabric Environment for Boot Mirroring	5
Using Prestoserve With VxVM	5
DMP Coexistence with EMC PowerPath	8
The DMP Command Utility (vxdmpadm)	10
Dynamic LUN Expansion	12
Installing an Array Support Library (ASL)	14
Commands Used to Maintain ASL Packages	14
Adding an ASL package	15
Removing an ASL Package	16
Upgrading an ASL Package	17
Partitioning Disks on Sun StorEdge 6120/6320 Arrays	17
ASLs Created on Previous Releases of VxVM	17
Extended Copy Service	18
Enabling a Disk for Ecopy Operation	18
Extended Copy Service Support for Hitachi 9900 and 9900V Arrays	19
JNI Host Bus Adapters (HBA)	20



VxVM Disk Array Configuration Updates	21
DMP and EMC Symmetrix Disk Array Support	21
HP SureStore™ E Disk Array XP256	22
HP SureStore™ E Disk Array XP512	24
IBM Enterprise Storage Server (ESS)	26
Support for IBM-SDD (Subsystem Device Driver)	29
Using Hitachi Arrays	29
DMP and Hitachi Data Systems SCSI Disk Array Subsystem Support	29
DMP and Hitachi Data Systems Fiber Disk Array Subsystem Support	30
HDS 5800E/9200	30
HDS 7700E/9910/9960/SE9960/SE9980	30
DMP and JBOD Disk Array Support	31
Discovering and Dynamically Adding Arrays	31
DMP Co-existence with Drivers	31
Using the AP Driver with VERITAS Volume Manager	32
Using the ATF Driver with VxVM	34
DMP and ATF Driver Co-existence Procedure	35
Sun StorEdge A5x00 Array	35
Dynamic Reconfiguration	37
Domain Configuration Type 1	38
Domain Configuration Type 2	39
Domain Configuration Type 3	41
Using the Sun StorEdge A3000/A3500 Array with VxVM	42
Installation of VxVM with the Sun StorEdge A3000/A3500 Array	43
VxVM and Sun StorEdge A3000/A3500 Array Device Nodes	43
Benefits in Using VxVM with the Sun StorEdge A3000/A3500 Array	44
Improved Manageability	44
Performance Analysis and Online Load Balancing Between LUNs	44
Increasing Capacity Online by LUN Concatenation	44
Online Data Movement Between LUNs	45



Stable Backup	45
Dynamic Partitioning	45
Sun StorEdge A3000/ A3500 Array Supported Configurations	45
Sun StorEdge A3000/ A3500 Array Unsupported Configurations	46
VxVM Hot-Relocation and Sun StorEdge A3000/ A3500 Array Hot Sparing	46
Using the Sun StorEdge T3/T3+/T4 Array with VxVM	46
Working with VxVM Disks on a Sun StorEdge T3 Array	47
Using a T3+ as a Boot Device	47
T3B firmware upgrade Considerations	48
Index	51





Preface

The *VERITAS Volume Manager Hardware Notes* provides hardware information on VERITAS Volume Manager (VxVM) Release 4.1.

This guide is intended for system administrators responsible for installing, configuring, and maintaining systems under the control of VxVM. The purpose of this manual is to provide the system administrator with information on hardware requirements for VxVM 4.1.

This guide assumes that you have a:

- ◆ working knowledge of the Solaris™ operating system
- ◆ basic understanding of Solaris system administration
- ◆ basic understanding of storage management

This manual describes how to use the VxVM command line interface for administering hardware with VxVM. Detailed descriptions of the VxVM commands and utilities, their options, and details on their use are located in the VxVM manual pages.

Note Most VxVM commands require superuser or other appropriate privileges.



Conventions

Convention	Usage	Example
monospace	Used for path names, commands, output, directory and file names, functions, and parameters.	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls(1)</code> manual page for more information.
monospace (bold)	Indicates user input.	# ls pubs C:\> dir pubs
<i>italic</i>	Identifies book titles, new terms, emphasized text, and variables replaced with a name or value.	See the <i>User's Guide</i> for details. The variable <i>system_name</i> indicates the system on which to enter the command.
bold	Depicts GUI objects, such as fields, list boxes, menu selections, etc. Also depicts GUI commands.	Enter your password in the Password field. Press Return .
blue text	Indicates hypertext links.	See " Getting Help " on page ix.
#	Unix superuser prompt (all shells).	# cp /pubs/4.0/user_book /release_mgnt/4.0/archive
C:\>	Windows user prompt.	C:\> copy \pubs\4.0\user_book c:\release_mgnt\4.0\archive

Getting Help

For technical assistance, visit <http://support.veritas.com> and select phone or email support. This site also provides access to resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and the VERITAS customer email notification service. Use the Knowledge Base Search feature to access additional product information, including current and past releases of product documentation.

Diagnostic tools are also available to assist in troubleshooting problems associated with the product. These tools are available on disc or can be downloaded from the VERITAS FTP site. See the `README.VRTSspt` file in the `/support` directory for details.

For license information, software updates and sales contacts, visit <https://my.veritas.com/productcenter/ContactVeritas.jsp>. For information on purchasing product documentation, visit <http://webstore.veritas.com>.

Documentation Feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to foundation_docs@veritas.com. Include the title and part number of the document (located in the lower left corner of the title page), and chapter and section titles of the text on which you are reporting. Our goal is to ensure customer satisfaction by providing effective, quality documentation. For assistance with topics other than documentation, visit <http://support.veritas.com>.





Hardware Notes

This document provides hardware support information for VERITAS Volume Manager (VxVM) Release 4.1.

The following topics are discussed in this guide:

- ◆ [Hardware Support in VxVM 4.1](#)
- ◆ [Arrays Supported in a Fabric Environment for Boot Mirroring](#)
- ◆ [Using Prestoserve With VxVM](#)
- ◆ [DMP Coexistence with EMC PowerPath](#)
- ◆ [The DMP Command Utility \(vxdkpadm\)](#)
- ◆ [Dynamic LUN Expansion](#)
- ◆ [Installing an Array Support Library \(ASL\)](#)
- ◆ [Extended Copy Service](#)
- ◆ [JNI Host Bus Adapters \(HBA\)](#)
- ◆ [VxVM Disk Array Configuration Updates](#)
- ◆ [Using the Sun StorEdge A3000/A3500 Array with VxVM](#)
- ◆ [Using the Sun StorEdge T3/T3+/T4 Array with VxVM](#)

Hardware Support in VxVM 4.1

This release of VxVM has been tested on:

- ◆ Solaris 8 (32- and 64-bit)
- ◆ Solaris 9 (32- and 64-bit)
- ◆ Solaris 10 (64-bit)

You can extract a list of supported hardware including arrays from the VERITAS customer support web site at <http://support.veritas.com>.



Disk Arrays Supported for DMP

DMP uses an Array Support Library (ASL) to provide support for specific vendors' disk arrays. The DMP feature of VxVM supports for the following arrays. If your array is not in this list, see “[Vendor ASL packages for Supported Arrays](#)” on page 3 for additional supported arrays.

- ◆ EMC Symmetrix™
- ◆ HP SureStore™ E Disk Array XP256/XP512/ XP48
- ◆ IBM Enterprise Storage Servers™ (ESS)
- ◆ Hitachi Data Systems 5800E/7700E/9910/9960/9200 Disk Array Subsystem
- ◆ Sun StorEdge™ 3900 and 6900 Series
- ◆ Sun StorEdge 3500, 6100, 6300 and 3300 series
- ◆ Sun StorEdge A5x00 Array
- ◆ Sun StorEdge T3 Array
- ◆ Sun StorEdge T3+ Array
- ◆ JBOD (Just a Bunch of Disks)
- ◆ SCSI 3 JBOD; arrays that support SCSI Inquiry Vital Product Data (VPD), pages 128 and 131.
- ◆ NEC arrays: DS1200, DS1200F, DS3000SL, DS3000SM, DS3001, DS3002, DS1000, DS1000F, DS1100, DS1100F, DS3011, DS1230, DS450, DS450F, iStorage 1000, iStorage 2000, iStorage 3000 and iStorage 4000.
- ◆ SEAGATE disks that return unique serial numbers in standard SCSI inquiry data.
- ◆ Storage Computer™ OmniRaid™ disk array. To multipath Storage Computer disk arrays connected to the system while using VxVM 4.1, you must assign a unique system name for each disk array connected to a machine. To find out how to set up a system name for storage computer disk arrays, see *The RAID-5 Users Manual* at the web site <http://www.storage.com>.
- ◆ ECCS™ Synchronix™ Array.
- ◆ Fujitsu GR710, GR720, GR730, GR740, GR820, and GR840.
- ◆ A full list of supported arrays can be found by entering “VERITAS Enabled” into the search engine at the VERITAS Support web site, <http://support.veritas.com>.

Note Arrays supported by the DMP feature are a subset of the arrays supported by VxVM. Other arrays that are not listed in this array list or in the next list can only be single path.

Vendor ASL packages for Supported Arrays

This section lists the disk arrays that are supported by Vendor ASL packages that can be installed as appropriate for your disk configuration. The following ASLs are supported with this release of VxVM; if your array is in this list, see [“Installing an Array Support Library \(ASL\)”](#) on page 14.

Vendor	Array models supported	arrays-asl ASL Package Name
3PARData	InServ Storage Servers	TPDasl
ADTX	FC-II	ADTXFCII
ADTX	AXRR-L	ADTXFCL
Chaparral	RIO	CNSiRIO
Chaparral	RIVA	CNSiRIVA
CISCO	Cisco MDS 9000	VRTScsco
Dothill	SanNet II FC	HLasl
Dothill	RIO Xtreme	DRIOX
EMC	Clariion FC4700, CX400,CX600	CLR-ASL
Fujitsu	E3000	FJSVasl
Fujitsu	E6000	FJSvasle6
Hitachi	9900v (9970/80),	HCL-ASL
Hitachi	9500v (9530/70/80)	HTC-9500V
Hitachi	9980V	HTC-ASL
HP	XP1024/128 with ECOPY Support	HPXP1
HP	VA7100/VA7110/ VA7400/VA7410	HPVA7xx0



Vendor	Array models supported	arrays-asl ASL Package Name
IBM	FastT Storage servers	SMibmasl
LSI	E series	SMlsiasl
Netapp	FAS9xx, F8xx, R100	NTAP-ASL
Storagetek	STK D178/280/bladeStor	SMstkasl
Sun	Sun StorEdge 3310	VRTS3310
Sun	Sun StorEdge 3510	VRTS3510
Sun	Sun StorEdge 3511	SUNW3511
Sun	Sun StorEdge 6120/6320	VRTSt4
Sun	Sun StorEdge 6130	SUNW6130
Sun	Sun StorEdge 6920	VRTS6920
Sun	Sun StorEdge 9980	HCL-ASL
Sun	Sun StorEdge 9900 (9970/9980)	HCL-ASL
Toshiba	AF1500/AF3500/AF2000/AF7000	TSBvxafsl
Toshiba	Arrayfort AF1500	TSBalaf

For the most up to date list of supported ASLs, you should visit The VERITAS support site at: <http://support.veritas.com>, and select "Volume Manager", "Volume Manager for UNIX", and click on "Compatibility". Enter the details (Product Version and Product Platform) into the box, and then select "HardWare Compatibility List". The Hardware Compatibility List includes a list of supported arrays, and a list of supported arrays together with a pointer to the appropriate Tech Note.

Note Arrays supported by the DMP feature are a subset of the arrays supported by VxVM. Other arrays that are not listed in this array list or in the previous list can only be single path.



Arrays Supported in a Fabric Environment for Boot Mirroring

For a list of supported arrays, go to <http://support.veritas.com>, and select “Volume Manager”, “Volume Manager for UNIX”, and click on “Compatibility”. You will be notified in Tech Notes as more arrays are tested and certified. You should also periodically visit the web site.

Using Prestoserve With VxVM

Prestoserve is designed to accelerate performance of file systems, particularly when used on a server for NFS advertised file systems. This is accomplished via the use of NVRAM hardware and the Prestoserve drivers. The hardware provides a fast, nonvolatile solid-state memory writeback cache that can cause writes to a disk device to be returned to the user as completed before the data reaches the disk.

This mechanism can be configured to work below VxVM as a direct replacement for the disk device that VxVM uses. This approach presents no particular problems for VxVM, which remains unaware of the underlying cache device. In the event of a failure of the NVRAM devices, it is possible to lose data since the disks backing the NVRAM may not be up-to-date.

Prestoserve can be configured to run above VxVM in such a way that VxVM replaces the disks that Prestoserve controls. In this situation, VxVM has a number of issues to address.

The most significant issue is with the startup of Prestoserve. Following a system failure, the Prestoserve drivers cause a flush of all outstanding dirty buffers to be flushed to disk. If this flush request occurs before VxVM drivers have been loaded into the kernel and before the volume devices can be started and made available for use, then Prestoserve’s attempts at flushing to the volumes fails.

Caution This problem can lead to file system corruption and data loss.

To prevent this situation, it is recommended that the order of the starting of Prestoserve with respect to the volumes be altered to occur after the volumes have been started. To achieve this, perform the following steps:

1. Edit the `/etc/system` file to add the following line:

```
exclude: drv/pr
```

2. Remove the `forceload` line:

```
forceload: drv/pr
```



This causes the Prestoserve driver to be loaded and the flush operation to be started *after* the volume devices have been started.

3. If your system is running Solaris 8 or Solaris 9, edit the file `/etc/init.d/vxvm-startup2` and add the following lines to the end of the file:

```
modload /kernel/drv/pr
presto -p > /dev/null
```

If your system is running Solaris 10, edit the file `/lib/svc/method/vxvm-startup2` and add the following lines to the end of the file:

```
modload /kernel/drv/pr
presto -p > /dev/null
```

This causes initialization of the Prestoserve driver (and flushing of saved blocks) following the start of all volumes.

Edit the `/etc/init.d/prestoserve` file and replace the following line:

```
presto -u
```

with the following line:

```
presto -u /mnt1 /mnt2 ...
```

Where `/mnt1` and `/mnt2` should list every file system to be accelerated with Prestoserve. Note that this list must *not* include `/`, `/usr`, `/usr/kvm`, `/var`, or `/var/adm`.

Caution Without the explicit starting of `presto` with the desired file systems, the default action accelerates all file systems, which (on the next reboot) can lead to data corruption. The file system checks for `/`, `/usr`, `/usr/kvm`, `/var`, and `/var/adm`, and they will precede the flushing of the Prestoserve buffers.

As is stated in the *Prestoserve User's Guide*, it is important not to reboot the system using the `-r` flag unless the system was cleanly shut down. VxVM can be used to add and remove drives without the need for a `boot -r`.

A second, more subtle, issue is in Prestoserve's use of disk devices. Some applications (including Prestoserve) maintain device numbers between reboots. VxVM also attempts to maintain device numbers between reboots. However, if a disk group is moved from one machine to another (or reattaching of disks causes some previously non-working disk group to reappear), it is possible for a conflict of minor numbers to be detected. VxVM

tries to prevent this by using randomly-assigned device number ranges (which are less likely to result in conflicts), but this policy can fail if two disk groups from different machines end up with the same randomly-assigned range.

The danger of VxVM changing its device numbers on a reboot following a system failure is that Prestoserve may flush its dirty buffers to the wrong volume devices. This may have destructive results.

If two disk groups do end up with volumes that use conflicting device numbers, the device number for one of the volumes is temporarily remapped on reboot (or when the second disk group is imported). Each reboot may choose a different remapping than was chosen on the previous reboot, leading to problems with Prestoserve. Any time a volume device number is temporarily remapped, the following console message displays:

```
VxVM vxconfigd WARNING V-5-1-663 Group ...: Duplicate virtual
device number:      Volume ... remapped from old-number to
new-number
```

If you see such a message, you should change the numbers for all volumes in the named disk group, using the `vxchg remminor` command.



DMP Coexistence with EMC PowerPath

With this release of VxVM, DMP can coexist with PowerPath. Depending on the scenario, you might need to install the EMC Cx600 ASL and its associated APM. To use DMP with PowerPath, you should be aware of the following scenarios.

- ◆ If you are installing VxVM 4.1 and PowerPath is installed, you do not need to install the Cx600 ASL and its associated APM. The array must be configured in explicit fail-over mode.
- ◆ If you are installing VxVM 4.1 and PowerPath is not installed, you must install the Cx600 ASL and its associated APM. The array can be in any mode.

The following table lists these scenarios and the required array modes:

PowerPath	DMP	Array mode
Installed	libvxemc.so handles EMC arrays and DGC claiming internally. PP handles fail-over. No need to install C.x600 ASL or Cx600 APM	Explicit fail-over
Not installed; the array is not Cx600	libvxemc.so handles the EMC Symmetrix array. DMP handles multipathing.	-
Not installed, but the array is Cx600	DMP handles multipathing. Cx600 ASL and Cx600 APM must be installed	-

Note If you are upgrading a system from VxVM 4.0 to VxVM 4.1, which has PowerPath installed, and the Cx600 ASL and its associated Cx600 APM are also installed, you must uninstall both the Cx600 ASL and APM, otherwise the Cx600 will claim the disks and the PowerPath discs will not be identified.

If you are upgrading a system from VxVM 4.0 to VxVM 4.1, which does not have



PowerPath installed, but the Cx600 ASL and its APM are both installed, then the Cx600 ASL and its associated APM should not be uninstalled.

If you are upgrading a system from VxVM 4.0 to VxVM 4.1, which has PowerPath installed, and the EMCpower discs are used as foreign discs, you must remove the foreign definitions using the `vxddladm rmforeign` command.

Foreign devices that were used in VxVM 4.0 are configured as `simple` disks with a default private region. To use them with VxVM 4.1, you must convert them to `auto:simple` disks by following the procedure described in the *Storage Foundation Installation Guide* under the heading *Converting Foreign Disks to auto:simple*.



The DMP Command Utility (vxddmpadm)

The `vxddmpadm` utility administers the DMP functionality of VxVM. Refer to the manual page of the `vxddmpadm(1M)` utility for a detailed description of the options and attributes of this command.

You can use the `vxddmpadm` utility option `listctlr` to list all controllers connected to disks that are attached to the host. For example, to list all controllers connected to disks on the host, use the following command:

```
# vxddmpadm listctlr all
```

to display the output from the above command:

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c0	OTHER	ENABLED	Disk
c1	Disk	ENABLED	Disk
c2	Disk	ENABLED	Disk

The previous displayed output shows that:

- ◆ The first controller `c0` is connected to disk(s) that are not under any recognized DMP category.
- ◆ The second and third controllers (`c1` and `c2`) are connected to a JBOD with SEAGATE disks.

All the above controllers are in the `ENABLED` state which indicates that they are available for I/O operations.

The state `DISABLED` is used to indicate that controllers are unavailable for I/O operations. The unavailability can be due to a hardware failure or due to I/O operations being disabled on that controller by the System Administrator. The controller state can be changed by using the `vxddmpadm` utility.

To list all the paths that are connected to a particular controller, you can use the `getsubpaths` option with the `ctlr` attribute. For example, use the following command:

```
# vxddmpadm getsubpaths ctlr=c1
```

to display the output from the above command:

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c1t0d0s2	ENABLED	-	c2t0d0s2	Disk	Disk0
c1t1d0s2	ENABLED	-	c2t1d0s2	Disk	Disk
c1t2d0s2	ENABLED	-	c2t2d0s2	Disk	Disk
c1t3d0s2	ENABLED	-	c2t3d0s2	Disk	Disk
c1t4d0s2	ENABLED	-	c2t4d0s2	Disk	Disk
c1t5d0s2	ENABLED	-	c2t5d0s2	Disk	Disk
c1t6d0s2	ENABLED	-	c2t6d0s2	Disk	Disk



c1t16d0s2	ENABLED	-	c2t16d0s2	Disk	Disk
c1t17d0s2	ENABLED	-	c2t17d0s2	Disk	Disk
c1t18d0s2	ENABLED	-	c2t18d0s2	Disk	Disk
c1t19d0s2	ENABLED	-	c2t19d0s2	Disk	Disk
c1t20d0s2	ENABLED	-	c2t20d0s2	Disk	Disk
c1t21d0s2	ENABLED	-	c2t21d0s2	Disk	Disk
c1t22d0s2	ENABLED	-	c2t22d0s2	Disk	Disk

The output display shows the paths that are connected to the controller named **c1**.

Additional information displays as follows:

- ◆ Path `c1t0d0s2` (represented by nodes in the `/dev/rdisk` and `/dev/dsk` directories) is in the **ENABLED** state.
- ◆ Path `c1t0d0s2` is represented by the DMP metanode `c2t0d0s2`, which is represented by device nodes in the `/dev/vx/dmp` and `/dev/vx/rdmp` directories.

You can use the `getsubpaths` option combined with the `dmpnodename` attribute to list all paths that are connected to a LUN (represented by a DMP device). For example, to list information about paths that lead to the LUN named `c1t0d0s2`, use the following command:

```
# vxddmpadm getsubpaths dmpnodename=c1t0d0s2
```

to display the output from the above command:

NAME	STATE	PATH-TYPE	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME
c2t0d0s2	DISABLED	-	c2	Disk	Disk
c1t0d0s2	ENABLED	-	c1	Disk	Disk

The listing above shows that the DMP device `c1t0d0s2` has two paths to it that are named `c1t0d0s2` and `c2t0d0s2`. Additional information indicates that only one of these paths is available for I/O operations. One of these paths is in the **ENABLED** state and the other is in the **DISABLED** state. Both paths are in a SEAGATE disk array.

To retrieve the name of the DMP device that controls a particular path, the `getdmpnode` option can be combined with the `nodename` attribute by using the following command:

```
# vxddmpadm getdmpnode nodename=c1t0d0s2
```

to display the output from the above command:

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
c2t0d0s2	ENABLED	Disk	2	2	0	Disk

This example output shows that the physical path `c1t0d0s2` is owned by the DMP device `c2t0d0s2`, which has two paths to it.



For information about using `vxddmpadm` with third part drivers, see the *Administering DMP chapter* in the *VERITAS Volume Manager Administrator's Guide*.

Dynamic LUN Expansion

The following form of the `vxdisk` command can be used to make VxVM aware of the new size of a virtual disk device that has been resized:

```
# vxdisk [-f] [-g diskgroup] resize {accessname|medianame} \  
  [length=value]
```

The device must have a SCSI interface that is presented by a smart switch, smart array or RAID controller. Following a resize operation to increase the length that is defined for a device, additional disk space on the device is available for allocation. You can optionally specify the new size by using the `length` attribute.

If a disk media name rather than a disk access name is specified, the disk group must either be specified using the `-g` option or the default disk group will be used. If the default disk group has not been set up, an error message will be generated.

Note A Storage Foundation license is required to use the `vxdisk resize` command.

This facility is provided to support dynamic LUN expansion by updating disk headers and other VxVM structures to match a new LUN size. It does not resize the LUN itself.

Any volumes on the device should only be grown after the device itself has first been grown. Otherwise, storage other than the device may be used to grow the volumes, or the volume resize may fail if no free storage is available.

Resizing should only be performed on devices that preserve data. Consult the array documentation to verify that data preservation is supported and has been qualified. The operation also requires that only storage at the end of the LUN is affected. Data at the beginning of the LUN must not be altered. No attempt is made to verify the validity of pre-existing data on the LUN. The operation should be performed on the host where the disk group is imported (or on the master node for a cluster-shared disk group).

Resizing of LUNs that are not part of a disk group is not supported. It is not possible to resize LUNs that are in the boot disk group (aliased as `boot.dg`), in a deported disk group, or that are offline, uninitialized, being reinitialized, or in an error state.

Caution Do not perform this operation when replacing a physical disk with a disk of a different size as data is not preserved.

Before reducing the size of a device, any volumes on the device should first be reduced in size or moved off the device. By default, the resize fails if any subdisks would be disabled as a result of their being removed in whole or in part during a shrink operation.

If the device that is being resized has the only valid configuration copy for a disk group, the `-f` option may be specified to forcibly resize the device.

Resizing a device that contains the only valid configuration copy for a disk group can result in data loss if a system crash occurs during the resize.

Resizing a virtual disk device is a non-transactional operation outside the control of VxVM. This means that the resize command may have to be re-issued following a system crash. In addition, a system crash may leave the private region on the device in an unusable state. If this occurs, the disk must be reinitialized, reattached to the disk group, and its data resynchronized or recovered from a backup.



Installing an Array Support Library (ASL)

VxVM provides Dynamic MultiPathing (DMP) support for new disk arrays in the form of Array Support Library (ASL) software packages. You can obtain ASL packages for new arrays from:

- ◆ The VxVM release CDROM CD1 in the `veritas_enabled` directory.
- ◆ The Storage Array Vendor's support site.
- ◆ The VERITAS Technical Support site, <http://support.veritas.com>. Select "Volume Manager" from the Product Family menu and "Volume Manager on Unix" from the Product men. Then click on Knowledge Base Search and enter "asl" followed by the vendor and/or model name.

VxVM 4.1 allows an ASL to be added to a running VxVM system (see "[Adding an ASL package](#)" on page 15). This means that you can add support for a particular disk array type (including multipathing support by Dynamic Multipathing, DMP) without having to: stop VxVM, reboot the system, or modify the VxVM package.

Note For the installed ASL to be discovered correctly, the array LUNs must be visible to the operating system through multiple paths, and they must be appropriately configured to support DMP. For instructions on how to configure the array correctly, please contact the array vendor.

You can also remove ASL packages from a running VxVM system (see "[Removing an ASL Package](#)" on page 16). If the associated disk array type contains active volumes, these will remain available during and after the addition or removal of the ASL package.

Only one version of an ASL for a specific type/model of disk array can be active in the system. To install a different version of an ASL, you need to first remove the previous version of the ASL package, then install the new package (see "[Upgrading an ASL Package](#)" on page 17).

Commands Used to Maintain ASL Packages

The following packaging commands are available for installing, removing, and listing ASL packages.

Install	Remove	Obtain Information
<code>pkgadd</code>	<code>pkgrm</code>	<code>pkginfo</code>

See the appropriate manual page for more information on these commands.



Adding an ASL package

Only one version of an ASL can be active in the system. If you find a version of the ASL is already installed, you need to refer to [Upgrading an ASL Package](#).

Note Disk arrays that are supported by an ASL do not have to be connected to the system when the ASL package is installed.

To find out if package (*ASL_name*) is already installed; enter:

```
# pkginfo -l ASL_name
```

You will find the ASL packages on the Storage foundation CDROM CD1 in the */mount_point/storage_foundation/veritas_enabled* directory. They have the format, *packagename.gz.tar*.

To add an ASL package from CD-ROM, proceed as follows:

1. Change directory to the ASL packages directory;

```
# cd /mount_point/storage_foundation/veritas_enabled
```

2. Copy the required package to a temporary directory:

```
# cp packagename.tar.gz /tmp
```

3. Change directory to the temporary directory:

```
# cd /tmp
```

4. Unzip and extract the package:

```
# gzcat packagename.tar.gz | tar -xvf -
```

5. Use `pkgadd` to install the package:

```
# pkgadd -d . ASL_Name
```

If `vxconfigd` is active, invoke Device Discovery Layer (DDL) to allow the newly installed ASL libraries to claim the disk arrays:

```
# vxdctl enable
```

If `vxconfigd` is not active, then invoke DDL the next time `vxconfigd` is activated.

(See VERITAS Volume Manager manual pages for command information.)



6. Verify that disk array was claimed by the *ASL_name* library:

```
# vxddmpadm listenclosure all
# vxddisk list cxtxdxs2
# vxddmpadm getsubpaths ctrl=cx
```

(where *x* is the appropriate control unit no.)

Note If disk array has multiple paths, use the above commands to verify that all paths to disk are claimed.

Removing an ASL Package

Before removing an ASL package, first find out if the package (*ASL_name*) is installed; enter:

```
# pkginfo -l ASL_name
```

1. To remove an ASL package you enter:

```
# pkgrm ASL_name
```

2. If *vxconfigd* is active, invoke Device Discovery Layer (DDL) through the *vxddctl* command to allow the arrays to be reclaimed:

```
# vxddctl enable
```

If *vxconfigd* is not active, then invoke DDL the next time *vxconfigd* is activated.

(See VERITAS Volume Manager manual pages for command information.)

3. Verify that disk array was claimed as "Disk".

```
# vxddmpadm listenclosure all
```

Note If the disk array has multiple paths, they must all be claimed as "DISK" (use the *vxddladm addjbod* command); otherwise, duplicate diskid errors will be detected.

Upgrading an ASL Package

Before upgrading an ASL package (installing a new version of an already installed package), first find out if the old package (*ASL_name*) is already present; enter:

```
# pkginfo -l ASL_name
```

1. Remove the package (*ASL_name*):

```
# pkgrm ASL_name
```

2. Proceed with [Adding an ASL package](#).

Partitioning Disks on Sun StorEdge 6120/6320 Arrays

Unless the correct installation command sequence is followed, partitioning of disks on array might fail, and you might see error messages of the form:

```
Volume manager: Partitioning T40_27 as a new disk.
VxVM vxdisk Error V-5-1-2913: Device T40_27: online failed:
Device path not valid
The partitioning of T40_27 failed.
```

This situation can be avoided by ensuring that you install the VxVM and Sun StorEdge 6120/6320 Array Support Library packages, in the following order:

1. Install VxVM - using `pkgadd` or `installvm`
2. Install VRTSt4 (ASL pkg) - using `pkgadd`
3. Set up VxVM using `vxinstall` and `vxdiskadm`

ASLs Created on Previous Releases of VxVM

If you have an ASL that was developed on an older release of VxVM, and is defined as Active/Passive (A/P) or Active/Passive with LUN group failover (A/PG), and you want to use multiple primary paths, you should change the array path selection policy to round-robin or balanced. For example, to change the I/O policy for Active/Passive array types to round-robin, use the following command:

```
# vxdmpadm setattr arraytype A/P iopolicy=round-robin
```

For further information, see the “Administering Dynamic Multipathing (DMP)” chapter in the *VERITAS Volume Manager Administrator’s Guide*.



Extended Copy Service

The Extended Copy Service feature of VxVM automatically works in tandem with the extended copy engines from array vendors. When VxVM detects that the source and destination devices are enabled for extended copy, VxVM automatically off loads copy requests to the array's copy manager.

The benefits of VxVM Extended Copy Service are:

- ◆ Non-disruptive copy operations from disk to disk. The host server remains online during the copy and the data being copied remains accessible to the server.
- ◆ Server-free copy operation. The copy operation is done between the array subsystem and the target disk. The data copy operation does not use any CPU or I/O resources on the host server.

To see whether the Extended Copy Service feature is enabled on a disk, use the `vxprint` command as shown in the following example. The feature is enabled if an `ecopy_enabled` entry appears in the `flags` line.

```
# vxprint -l disk03
Disk group: rootdg

Disk:      disk03
info:      diskid=1234567890.59.vm250e1.veritas.com
assoc:     device=c2t2d0s2 type=auto
flags:     autoconfig ecopy_enabled
device:    path=/dev/vx/dmp/c2t2d0s4
devinfo:   publen=35354136 privlen=9167
```

Note If required, you can use the `-o noecopy` option to turn off Extended Copy Service for each invocation of the `vxplex att`, `cp`, `mv` and `snapstart` commands, and the `vxsd mv` command.

Enabling a Disk for Ecopy Operation

To enable a disk for ecopy operation:

1. Install the Hardware Assisted copy license
2. Enable the Extended copy features in the array. (This procedure is vendor specific.)
3. Install the vendor ASL that supports the ecopy feature. contact VITA@veritas.com for vendor ASL information.

Extended Copy Service Support for Hitachi 9900 and 9900V Arrays

To implement extended copy for the Hitachi 9900 and 9900V arrays, you must create the two files, `user_pwwn_file` and `user_luid_file`, that identify the Hitachi disk identification information. To create these files, run the `vxwwnluid` utility that is located in `/etc/vx/diag.d`. There are no parameters for the `vxwwnluid` utility, and the command must be executed as root.

- ◆ `/etc/vx/user_pwwn_file`

This file contains the disk access name and the *port world-wide-name* (pwwn) for each disk in the array. For the hitachi arrays, both the source and the destination devices must have entries in the this file. The information for each disk in the array is defined on a single line. The disk access name and PWWN are separated by a single tab character.

The following are sample entries from the Solaris `/etc/vx/user_pwwn_file` file:

```
sde      50060e800404040b
sdf      50060e800404040b
sdg      50060e800404040b
```

- ◆ `/etc/vx/user_luid_file`

This file contains the disk access names and their corresponding LUN numbers in the array. The information for each disk in the array is defined on a single line. The disk access name and the LUN are separated by a single tab character.

The following are sample entries from the `/etc/vx/user_luid_file` file:

```
sde      1
sdf      2
sdg      1
```



JNI Host Bus Adapters (HBA)

If your JNI card has Model numbers FCE-1063, FCE2-1063, FCE-6410, FCE2-6410, or FCE2-6412, then you may experience error messages of the form:

```
"Oct 22 00:16:16 ds13un jnic: [ID 847178 kern.notice] jnic1: Memory
port parity error detected
Oct 22 00:16:16 ds13un jnic: [ID 229844 kern.notice] jnic1: Link
Down
Oct 22 00:16:16 ds13un jnic: [ID 744007 kern.notice] jnic1:
Target0: Port
0000EF (WWN 500060E802778702:500060E802778702) offline.
Oct 22 00:16:18 ds13un jnic: [ID 709123 kern.notice] jnic1: Link Up
Oct 22 00:16:18 ds13un jnic: [ID 236572 kern.notice] jnic1:
Target0: Port
0000EF (WWN 500060E802778702:500060E802778702) online.
Oct 22 00:16:18 ds13un jnic: [ID 229844 kern.notice] jni
Contact JNI support for more information."
```

As a workaround, you should add the following parameter to the JNI configuration file (`jnic.conf`):

```
FcEnableContextSwitch = 1;
```



VxVM Disk Array Configuration Updates

This section lists some of the disk arrays supported for Dynamic Multipathing (DMP) and describes how to use VxVM with the Sun StorEdge A3000/A3500 Array.

DMP and EMC Symmetrix Disk Array Support

DMP works with the EMC Symmetrix disk array only if the disk array is configured in the *Common Serial Number Mode*. This configuration provides unique device identification on different paths for DMP to configure its database. Without this mode, DMP cannot provide multipathing functionality. DMP cannot access EMC Symmetrix disks as individual disks with a single path it found by scanning the device information tree in the kernel.

For the DMP functionality to work with the EMC Symmetrix disk array, the following conditions must be met:

- ◆ It is necessary that the EMC Symmetrix users program their disk arrays in the *Common Serial Number Mode* before installing the VERITAS Volume Manager package.
- ◆ The minimum level of microcode on the EMC Symmetrix must be as follows:

Sym-4	3330, 3430, 3700	5265 and 5266 microcode
Sym-4.8	3630, 3830, 3930	5265 and 5266 microcode
Sym-5	8130, 8430, 8730	5566 microcode

- ◆ The *Common Serial Number Mode* should be enabled for all channels that participate in DMP. This is done at installation time in the Symmetrix configuration by the EMC Customer Engineer.
- ◆ In addition, the *disable queue reset on unit attention mode* must be set for all channels that participate in the DMP. This also is done at installation time in the Symmetrix configuration by the EMC Customer Engineer.

Contact your EMC Technical Support Representative for instructions.

DMP supports SCSI and SAN data gateways using JNI and QLOGIC cards.

DMP supports the Ultra SCSI and the fast-wide differential (FWD) SCSI controllers used with the EMC Symmetrix disk array.

- ◆ Ultra SCSI and fast-wide differential (FWD) SCSI support has been tested on the EMC OSD (Open System Director) controller.



- ◆ DMP was tested on and supports the Jaycor fiber controllers: firmware release level F, and driver level V2.0.1EMC and V2.0.2EMC. You should contact Jaycor and EMC for the latest supported driver for the specific controller being used on the EMC Symmetrix.

Note Uncommenting the *failover=* line in */kernel/drv/fca.conf* is important for the DMP failover as it reduces the amount of time it takes for a path to be considered failed. The default time is over five minutes. Setting this value lower reduces the time considerably.

Contact your VERITAS sales representative for further updates on hardware support.

HP SureStore™ E Disk Array XP256

The minimum software driver and firmware levels for the JNI adapter are:

- ◆ FC64-1063 Software Driver: Minimum level is V2.3. Current level V2.4.1 is recommended.
- ◆ FC64-1063 Firmware level: Minimum level 13.3.5, or above.

Output examples include:

```
# vxdmpadm listctlr all
```

CTLR-NAME	ENCLR-TYPE	STATE	E NCLR-NAME
c0	Disk	ENABLED	Disk
c5	XP256	ENABLED	xp 2560
c6	XP256	DISABLED	xp 2560

```
# vxdmpadm getsubpaths ctlr=c5
```

NAME	STATE	PATH-TYPE (M)	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
c5t11d1s2	ENABLED	-	c5t11d1s2	XP256	xp2560	-
c5t11d2s2	ENABLED	-	c5t11d2s2	XP256	xp2560	-
c5t11d3s2	ENABLED	-	c5t11d3s2	XP256	xp2560	-
c5t11d4s2	ENABLED	-	c5t11d4s2	XP256	xp2560	-
c5t11d5s2	ENABLED	-	c5t11d5s2	XP256	xp2560	-
c5t11d6s2	ENABLED	-	c5t11d6s2	XP256	xp2560	-
c5t11d7s2	ENABLED	-	c5t11d7s2	XP256	xp2560	-
c5t11d8s2	ENABLED	-	c5t11d8s2	XP256	xp2560	-
c5t11d9s2	ENABLED	-	c5t11d9s2	XP256	xp2560	-
c5t11d10s2	ENABLED	-	c5t11d10s2	XP256	xp2560	-
c5t11d11s2	ENABLED	-	c5t11d11s2	XP256	xp2560	-
c5t11d12s2	ENABLED	-	c5t11d12s2	XP256	xp2560	-

vxdisk list

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:sliced	-	-	online invalid
c2t8d0s2	auto:sliced	-	-	online
c2t9d0s2	auto:sliced	-	-	online
c2t11d0s2	auto:sliced	-	-	online
c2t12d0s2	auto:sliced	-	-	online
c5t13d0s2	auto:sliced	-	-	online
c8t0d0s2	auto:sliced	disk01	testdg	online
c8t1d0s2	auto:sliced	disk02	testdg	online
c8t2d0s2	auto:sliced	disk03	testdg	online
c8t3d0s2	auto:sliced	disk04	testdg	online
c6t4d0s2	auto:sliced	disk05	testdg	online

vxdisk list disk03

```

Device:      c8t2d0s2
devicetag:  c8t2d0
type:       auto
hostid:     vm250e1.veritas.com
disk:       name=disk03 id=961801990.1071.vm250e1.veritas.com
group:      name=testdg id=961801968.1025.vm250e1.veritas.com
info:       format=sliced, privoffset=1, pubslice=4, privslice=3
flags:      online ready private autoconfig autoimport imported
pubpaths:  block=/dev/vx/dmp/c8t2d0s4
            char=/dev/vx/rdmp/c8t2d0s4
privpaths: block=/dev/vx/dmp/c8t2d0s3
            char=/dev/vx/rdmp/c8t2d0s3
version:    2.1
iosize:     min=512 (bytes) max=2048 (blocks)
public:     slice=4 offset=0 len=14417280 disk_offset=5112
private:    slice=3 offset=1 len=2879 disk_offset=2556
update:     time=1063149836 seqno=0.46
ssb:        actual_seqno=0.0
headers:    0 24 8
configs:    count=1 len=1856
logs:       count=1 len=281
Defined regions:
c onfig    priv 000017-000247[000231]: copy=01 offset=000000
           enabled
c onfig    priv 000249-002121[001873]: copy=01 offset=000231
           enabled
l og       priv 002122-002439[000318]: copy=01 offset=000000
           enabled
Multipathing information:
numpaths:  2
c5t11d12s2    state=enabled
c6t8d12s2    state=enabled

```



HP SureStore™ E Disk Array XP512

The minimum software driver and firmware levels for the JNI adapter are:

- ◆ FC64-1063 Software Driver: Minimum level is V2.3. Current level V2.4.1 is recommended.
- ◆ FC64-1063 Firmware level: Minimum level 13.3.5, or above.

Note When using an XP512 array, and the array is attached through a JNI HBA(FCI-1063), you should use the Solaris-Hitachi fca-pci package which can be found at <http://www.jni.com/Drivers>. When you reach this page, you should select FCI-1063 and then download and install the package, SOLARIS_PCI_258HIT0701.tar

Output examples include:

```
# vxddmpadm listctlr all
```

CTLR-NAME	ENCLR-TYPE	STATE	E NCLR-NAME
c0	SEAGATE	ENABLED	se agate0
c2	SEAGATE	ENABLED	se agate0
c5	XP512	ENABLED	xp 5120
c6	XP512	ENABLED	xp 5120

```
# vxddmpadm getsubpaths ctlr=c5
```

NAME	STATE	PATH-TYPE (M)	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
c5t0d0s2	ENABLED	-	c5t0d0s2	XP512	xp5120	-
c5t0d1s2	ENABLED	-	c5t0d1s2	XP512	xp5120	-
c5t0d2s2	ENABLED	-	c5t0d2s2	XP512	xp5120	-
c5t0d3s2	ENABLED	-	c5t0d3s2	XP512	xp5120	-
c5t0d4s2	ENABLED	-	c5t0d4s2	XP512	xp5120	-
c5t0d5s2	ENABLED	-	c5t0d5s2	XP512	xp5120	-
c5t0d6s2	ENABLED	-	c5t0d6s2	XP512	xp5120	-
c5t0d7s2	ENABLED	-	c5t0d7s2	XP512	xp5120	-
c5t0d9s2	ENABLED	-	c5t0d9s2	XP512	xp5120	-
c5t0d10s2	ENABLED	-	c5t0d10s2	XP512	xp5120	-
c5t0d11s2	ENABLED	-	c5t0d11s2	XP512	xp5120	-
c5t0d12s2	ENABLED	-	c5t0d12s2	XP512	xp5120	-
c5t0d13s2	ENABLED	-	c5t0d13s2	XP512	xp5120	-
c5t0d14s2	ENABLED	-	c5t0d14s2	XP512	xp5120	-
c5t0d15s2	ENABLED	-	c5t0d15s2	XP512	xp5120	-

vxddmpadm getsubpaths ctrl=c6

NAME	S TATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	E NCLR-NAME
c6t0d0s2	ENABLED	-	c 5t0d0s2	X P512	xp5120
c6t0d1s2	ENABLED	-	c 5t0d1s2	X P512	xp5120
c6t0d2s2	ENABLED	-	c 5t0d2s2	X P512	xp5120
c6t0d3s2	ENABLED	-	c 5t0d3s2	X P512	xp5120
c6t0d4s2	ENABLED	-	c 5t0d4s2	X P512	xp5120
c6t0d5s2	ENABLED	-	c 5t0d5s2	X P512	xp5120
c6t0d6s2	ENABLED	-	c 5t0d6s2	X P512	xp5120
c6t0d7s2	ENABLED	-	c 5t0d7s2	X P512	xp5120
c6t0d9s2	ENABLED	-	c 5t0d9s2	X P512	xp5120
c6t0d10s2	ENABLED	-	c 5t0d10s2	X P512	xp5120
c6t0d11s2	ENABLED	-	c 5t0d11s2	X P512	xp5120
c6t0d12s2	ENABLED	-	c 5t0d12s2	X P512	xp5120
c6t0d13s2	ENABLED	-	c 5t0d13s2	X P512	xp5120
c6t0d14s2	ENABLED	-	c 5t0d14s2	X P512	xp5120
c6t0d15s2	ENABLED	-	c 5t0d15s2	X P512	xp5120

vxddisk list

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:sliced	-	-	online invalid
c2t8d0s2	auto:sliced	-	-	online
c2t9d0s2	auto:sliced	-	-	online
c2t11d0s2	auto:sliced	-	-	online
c2t12d0s2	auto:sliced	-	-	online
c5t13d0s2	auto:sliced	-	-	online
c8t0d0s2	auto:sliced	disk01	testdg	online
c8t1d0s2	auto:sliced	disk02	testdg	online
c8t2d0s2	auto:sliced	disk03	testdg	online
c8t3d0s2	auto:sliced	disk04	testdg	online
c6t4d0s2	auto:sliced	disk05	testdg	online

vxddisk list c5t0d1s2

```

Device:      c5t0d1s2
devicetag:  c5t0d1
type:       auto
hostid:     vm250e1.veritas.com
disk:       name=disk02 id=976924352.1097.vm250e1.veritas.com
group:      name=testdg id=976924305.1025.vm250e1.veritas.com
info:       format=sliced, privoffset=1, pubslice=4, privslice=3
flags:      online ready private autoconfig autoimport imported
pubpaths:   block=/dev/vx/dmp/c5t0d1s4 char=/dev/vx/rdmp/c5t0d1s4
privpaths:  block=/dev/vx/dmp/c5t0d1s3 char=/dev/vx/rdmp/c5t0d1s3
version:    2.1
iosize:     min=512 (bytes) max=2048 (blocks)

```



```
public:  slice=4 offset=0 len=14417280 disk_offset=5112
private: slice=3 offset=1 len=2879 disk_offset=2556
update:  time=976952456 seqno=0.88
ssb:     actual_seqno=0.0
headers: 0 24 8
configs: count=1 len=2104
logs:    count=1 len=318
Defined regions:
con fig  priv 000017-000247[000231]: copy=01 offset=000000 enabled
con fig  priv 000249-002121[001873]: copy=01 offset=000231 enabled
log      pr iv 002122-002439[000318]: copy=01 offset=000000 enabled
Multipathing information:
numpaths: 2
c5t0d1s2      state=enabled
c6t0d1s2      state=enabled
```

IBM Enterprise Storage Server (ESS)

Minimum software driver and firmware levels for the JNI adapter are:

- ◆ FC64-1063 Software Driver: Minimum level is V2.3. Current level V2.4.1 is recommended.
- ◆ FC64-1063 Firmware level: Minimum level 13.3.5, or above.

Required Solaris Operating System parameters (in file /etc/system) for JNI HBA are:

- ◆ `sd:sd_max_throttle`: Determine the value by using the following formula:
 $256 / \text{LUNs per adapter}$
where LUNs per adapter is the largest number of LUNs assigned to a single adapter.
- ◆ `set sd:sd_io_time=0x78`
- ◆ `set sd:sd_retry_count=5`

Note For detailed information, please refer to the IBM ESS Host Attachment Guide.
This document can be found online at:
www.storage.ibm.com/hardsoft/products/ess/support/hostattach.pdf

Output examples include:

```
# vxdmpadm listctlr all
```

CTLR-NAME	ENCLR-TYPE	STATE	E NCLR-NAME
c0	OTHER	ENABLED	ot hers0
c3	IBM-SHARK	ENABLED	IB M_SHARK0
c4	IBM-SHARK	DISABLED	IB M_SHARK0
c0	SEAGATE	ENABLED	se agate0

```
# vxdmpadm getsubpaths ctlr=c4
```

NAME	STATE	PATH-TYPE (M)	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
c4t1d2s2	ENABLED	-	c3t0d2s2	IBM-SHARK	IBM_SHARK0	-
c4t1d3s2	ENABLED	-	c4t1d3s2	IBM-SHARK	IBM_SHARK0	-
c4t1d4s2	ENABLED	-	c3t0d4s2	IBM-SHARK	IBM_SHARK0	-
c4t1d5s2	ENABLED	-	c3t0d5s2	IBM-SHARK	IBM_SHARK0	-
c4t1d6s2	ENABLED	-	c3t0d6s2	IBM-SHARK	IBM_SHARK0	-
c4t1d7s2	ENABLED	-	c3t0d7s2	IBM-SHARK	IBM_SHARK0	-
c4t1d8s2	ENABLED	-	c3t0d8s2	IBM-SHARK	IBM_SHARK0	-
c4t1d9s2	ENABLED	-	c3t0d9s2	IBM-SHARK	IBM_SHARK0	-
c4t1d10s2	ENABLED	-	c3t0d10s2	IBM-SHARK	IBM_SHARK0	-
c4t1d11s2	ENABLED	-	c3t0d11s2	IBM-SHARK	IBM_SHARK0	-
c4t1d12s2	ENABLED	-	c3t0d12s2	IBM-SHARK	IBM_SHARK0	-
c4t1d13s2	ENABLED	-	c3t0d13s2	IBM-SHARK	IBM_SHARK0	-
c4t1d14s2	ENABLED	-	c3t0d14s2	IBM-SHARK	IBM_SHARK0	-

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:sliced	-	-	online invalid
c2t8d0s2	auto:sliced	-	-	online
c2t9d0s2	auto:sliced	-	-	online
c2t11d0s2	auto:sliced	-	-	online
c2t12d0s2	auto:sliced	-	-	online
c5t13d0s2	auto:sliced	-	-	online
c8t0d0s2	auto:sliced	disk01	testdg	online
c8t1d0s2	auto:sliced	disk02	testdg	online
c8t2d0s2	auto:sliced	disk03	testdg	online
c8t3d0s2	auto:sliced	disk04	testdg	online
c6t4d0s2	auto:sliced	disk05	testdg	online

```
# vxdisk list disk08
```

```
Device:      c4t1d3s2
devicetag:  c4t1d3
type:       auto
hostid:     vm250e1.veritas.com
disk:       name=disk08 id=963015532.1097.vm250e1.veritas.com
group:      name=testdg id=963015492.1025.vm250e1.veritas.com
info:       format=sliced, privoffset=1, pubslice=4, privslice=3
flags:      online ready private autoconfig autoimport imported
```



```

pubpaths:  block=/dev/vx/dmp/c4t1d3s4  char=/dev/vx/rdmp/c4t1d3s4
privpaths: block=/dev/vx/dmp/c4t1d3s3  char=/dev/vx/rdmp/c4t1d3s3
version:    2.1
iosize:    min=512 (bytes) max=256 (blocks)
public:    slice=4 offset=0 len=3897600 disk_offset=5112
private:   slice=3 offset=1 len=3839 disk_offset=2556
update:    time=963517507 seqno=0.213
ssb:       actual_seqno=0.0
headers:   0 24 8
configs:   count=1 len=2808
logs:      count=1 len=425
Defined regions:
con fig   priv 000017-000247[000231]: copy=01 offset=000000 enabled
con fig   priv 000249-002825[002577]: copy=01 offset=000231 enabled
log       pr iv 002826-003250[000425]: copy=01 offset=000000 enabled
Multipathing information:
numpaths: 2
c4t1d3s2   state=disabled
c3t0d3s2   state=enabled

```

vxddmpadm listctlr all

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c0	OTHER	ENABLED	ot hers0
c3	IBM-SHARK	ENABLED	IB M_SHARK0
c4	IBM-SHARK	DISABLED	IB M_SHARK0
c0	SEAGATE	ENABLED	se agate0

vxddmpadm getsubpaths ctlr=c3

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c3t0d1s2	ENABLED	-	c 3t0d1s2	IBM-SHARK	I BM_SHARK0
c3t0d2s2	ENABLED	-	c 3t0d2s2	IBM-SHARK	I BM_SHARK0
c3t0d4s2	ENABLED	-	c 3t0d4s2	IBM-SHARK	I BM_SHARK0
c3t0d5s2	ENABLED	-	c 3t0d5s2	IBM-SHARK	I BM_SHARK0
c3t0d6s2	ENABLED	-	c 3t0d6s2	IBM-SHARK	I BM_SHARK0
c3t0d7s2	ENABLED	-	c 3t0d7s2	IBM-SHARK	I BM_SHARK0
c3t0d8s2	ENABLED	-	c 3t0d8s2	IBM-SHARK	I BM_SHARK0
c3t0d9s2	ENABLED	-	c 3t0d9s2	IBM-SHARK	I BM_SHARK0
c3t0d10s2	ENABLED	-	c 3t0d10s2	IBM-SHARK	I BM_SHARK0
c3t0d11s2	ENABLED	-	c 3t0d11s2	IBM-SHARK	I BM_SHARK0
c3t0d12s2	ENABLED	-	c 3t0d12s2	IBM-SHARK	I BM_SHARK0
c3t0d13s2	ENABLED	-	c 3t0d13s2	IBM-SHARK	I BM_SHARK0
c3t0d14s2	ENABLED	-	c 3t0d14s2	IBM-SHARK	I BM_SHARK0
c3t0d3s2	ENABLED	-	c 4t1d3s2	IBM-SHARK	I BM_SHARK0



Support for IBM-SDD (Subsystem Device Driver)

VxVM now supports the co-existence of IBM-SDD (vpath) on the Solaris Operating System; the supported version is IBM-SDD Version 1.4.0.7. SDD supports the multipath configuration environments in the IBM ESS. VxVM uses enclosure-based names for SDD devices (vpath).

Support for SDD is provided by the new ASL, `libvxvpath.so`. vpath devices/metanodes are represented as `vpathN`, for example, `/dev/rdisk/vpath12[a-h]`, where *a-h* represents partition number or slices.

Using Hitachi Arrays

If you are considering having your host disk reside on an hitachi array, you should first contact Hitachi Data Systems for the latest information on supported boot disk configurations with hitachi arrays,

DMP and Hitachi Data Systems SCSI Disk Array Subsystem Support

DMP supports the Hitachi Data Systems 5700E and 5800 Disk Array Subsystems under the following conditions:

- ◆ The optional second controller is installed in the array.
- ◆ One controller is configured as *Active* and the other controller is configured as *Passive* (for improved performance).
- ◆ The array is connected to two SCSI cards in the host machine.
- ◆ The array has a firmware version that supports serial numbers in the SCSI inquiry data and they are configured correctly.

Supported Microcode Levels

Unit	Microcode Version	Flash Revision
DF350 5700E	x308/E	pb28
DF350F 5700E	x356/F	pe27
DF300 5700	x109/L	dq3x/4x (x=89)
DF400 5800	x407/M	a03, b13



Supported Microcode Levels

Unit	Microcode Version	Flash Revision
DF400 5800	x457/H	a03, b13

You can extract a list of supported hardware including arrays from the VERITAS customer support web site at: <http://support.veritas.com>

For the Hitachi DF array families (DF350, DF400, DF500), only A/P mode will be supported. If you have any of the Hitachi DF arrays installed, you must ensure that the arrays are configured as A/P mode.

Contact Veritas support for assistance in determining your Hitachi array mode.

Contact your Hitachi Data Systems Technical Support Representative for assistance in reconfiguring the DP arrays to A/P mode.

DMP and Hitachi Data Systems Fiber Disk Array Subsystem Support

HDS 5800E/9200

DMP supports HDS 5800E and 9200 Disk Array Subsystems in Active/Passive mode only.

HDS 7700E/9910/9960/SE9960/SE9980

DMP supports the HDS 7700E, 9910, 9960, SE9960 and SE9980 Disk Array Subsystems in Active/Active mode only.

For HITACHI model 9960, depending on the workload of the system, DMP performance can be improved by using the `vxddmpadm` command to adjust the partition size for the balanced I/O policy to a value higher than the system default of 2048 blocks. The recommended value is in the range of 32,768 to 131,072 blocks. See the *VERITAS Volume Manager Administrator's Guide* for details.

A full list of supported arrays can be found by going to the VERITAS Support web site, <http://support.veritas.com> and entering "VERITAS Enabled" into the search engine.

DMP and JBOD Disk Array Support

DMP on JBOD drives enables DMP to multipath JBOD disk arrays connected to a system using VxVM. These JBOD disk arrays/disks can contain disks from any vendor. However, DMP can correctly detect multiple paths to these disks/disk arrays only if the following conditions are satisfied:

- ◆ The serial number field of the Standard SCSI INQUIRY DATA (bytes 36—47) should be “world wide unique” for any disk that is connected to a JBOD. This fact must be known to the user before the user attempts to configure DMP to add these disks to the JBOD category of DMP.

Caution If any disk does not provide a unique serial number, DMP cannot detect multiple paths to the disk correctly and results are unpredictable. Contact the JBOD disk/disk array vendors to get correct information about unique serial numbers on disks.

- ◆ JBOD disk arrays that require multipathing functionality from DMP should be of the Active/Active type.

Once it is known that disks connected to a JBOD disk array satisfy the above conditions, you must use the `vxddladm addjbod` command to add disks to the JBOD category. See `vxddladm (1)` for further information.

Discovering and Dynamically Adding Arrays

You can dynamically add support for a new type of disk array which has been developed by a third party vendor. The support comes in the form of vendor-supplied libraries and is added to the Solaris Operating system by using the `pkgadd` command. For further information, see the *VERITAS Volume Manager Administrator's Guide*.

DMP Co-existence with Drivers

DMP supports co-existence with the following arrays:

- ◆ Sun's Alternate Pathing driver version 2.3.1 along with Solaris Operating system patch 110722-01 (or higher) for A5x00 and Sun StorEdge T3 disk arrays.
- ◆ DG CLARiiON with the ATF driver installed on the system.
- ◆ SYMBIOS Sun StorEdge A3000/A3500 Array, only when the RDAC driver is installed on the system.
- ◆ IBM-SDD (vpath)
- ◆ MPxIO including Sun StorEdge Traffic Manager



Using the AP Driver with VERITAS Volume Manager

This release of VERITAS Volume Manager co-exists only with version 2.3.1 of Sun's Alternate Pathing (AP) solution along with the Solaris Operating System patch 110722-01 (or higher). If you use AP, upgrade AP to 2.3.1 and install the Solaris Operating System patch 110722-01 before installing/upgrading to this release of VERITAS Volume Manager.

Note The `vxvmp` driver must always be present on the system for VxVM to function.

If the system has a version of AP lower than 2.3.1 or the system does not have the Solaris Operating System patch 110722-01, the installation/upgrade to VxVM 4.1 will abort with the following messages.

If the system has a version of AP lower than 2.3.1, the following message is displayed.

```
Alternate Pathing (AP) version 2.3.0 has been detected.
This version of VxVM co-exists only with AP 2.3.1 along with
following patch(es):
    110722-01
Please upgrade your AP package and install the above
patch(es) before you continue the VxVM installation. This
installation will now be aborted.
```

If the system has AP version 2.3.1 but the Solaris Operating System patch 110722-01 is not installed, the following message is displayed.

```
Alternate Pathing (AP) version 2.3.1 has been detected.
This version of VxVM requires the following patch(es) to co-exist
with AP 2.3.1:
    110722-01
Install the above patch(es) before you continue the VxVM
installation. This installation will now be aborted.
```

An upgrade to VERITAS Volume Manager 4.1 enables DMP if it was disabled earlier.

If there are devices under AP control then they will be claimed by DMP in the AP category. In the following example, an A5x00 device that is under AP control is claimed in the AP category.

```
# vxvmpadm listctlr all
```

```
CTLR-NAME          ENCLR-TYPE          STATE          E NCLR-NAME
=====
c0                  OTHER               ENABLED        ot hers0
mc2                 AP_SENA             ENABLED        AP_SENA0
```

```

# vxddmpadm getsubpaths ctrl=mc2
NAME          STATE      PATH-TYPE (M)  DMPNODENAME  ENCLR-TYPE  ENCLR-NAME  ATTRS
=====
mc2t35d0s2    ENABLED    -              mc2t35d0s2   AP_SENA     AP_SENA0    -
mc2t33d0s2    ENABLED    -              mc2t33d0s2   AP_SENA     AP_SENA0    -
mc2t34d0s2    ENABLED    -              mc2t34d0s2   AP_SENA     AP_SENA0    -

# vxddmpadm listenclosure all
ENCLR_NAME      EN CLR_TYPE    ENCLR_SNO      S  STATUS
=====
others0         OTHER          OTHER_DISKS    C  CONNECTED
SENA0           SENA          5080020000040750  D  DISCONNECTED
AP_SENA0        AP_SENA       5080020000040750  C  CONNECTED

# vxddisk list
DEVICE          TYPE           DISK           GROUP          STATUS
=====
c0t0d0s2       suto:sliced   -              -              offline
c0t1d0s2       auto:sliced   disk02         tes tdg        online
mc2t33d0s2     auto:sliced   -              -              offline
mc2t34d0s2     auto:sliced   disk03         tes tdg        online
mc2t35d0s2     auto:sliced   disk01         tes tdg        online

```

In the above example, *SENA0* is the enclosure under which the devices were categorized before the devices were added to AP control. The name assigned to the enclosure seen through AP is *AP_SENA0*.

If a pathgroup is added under AP control after VxVM is installed on a system, the `vxddctl enable` command must be executed for the AP metanodes to be visible to VxVM. After this is done, all devices in the disk array controlled by AP will be claimed in the AP category.

If a pathgroup under AP control is to be removed, the following procedure should be used.

1. All disks that are in this pathgroup must be removed entirely from VxVM control. Execute `vxddiskadm` and choose the `Remove a disk for replacement` option. Use this option to remove all AP metanodes from VxVM.
2. Run the following command on all AP metanodes.

```
# vxddisk rm daname
```

 where `daname` is the `c#t#d#` name
3. Follow the steps in the AP documentation to remove a pathgroup.



4. Run the VxVM configuration daemon, as follows:

```
# vxctl enable
```
5. Re-add the disks to VxVM by entering the command `vxdiskadm` and choose the following option:

```
Replace a failed disk or removed disk
```
6. Restart all volumes, if any, that were stopped as a result of the above operations.

Using the ATF Driver with VxVM

CLARiiON ATF driver and DMP co-existence is supported in VM 3.0 and above.

For DMP to co-exist with the ATF driver, the `CLatf` package (Application Transparent Failover Driver for CLARiiON) must be installed before the `VRTSvxvm` or `SUNWvxvm` packages are installed on the system.

Note In VxVM 4.1, ATF/DMP co-existence works *only* with the `CLatf` 1.3.1, 3.1.0, 3.3.0.3 and 3.4.0.7 packages.

For DMP to co-exist with the ATF driver, the boot disk should not be under control of the ATF driver.

You can make DMP co-exist with the ATF driver. To do this, use the procedure that is appropriate to your situation.

▼ **If VxVM is not installed, do the following:**

1. Install ATF.
2. Install VxVM.

▼ **If VxVM is installed with DMP currently enabled, do the following:**

1. Prevent or suppress DMP.
2. Install the ATF driver.
3. Install VxVM 4.1, if appropriate.
4. If DMP is not re-enabled, re-enable it now.

▼ **If VxVM is installed with DMP currently prevented or suppressed, do the following:**

1. Install the ATF driver.
2. Upgrade VxVM 4.1, if appropriate.
3. If necessary, reinstate the DMP settings.

DMP and ATF Driver Co-existence Procedure

The following warning messages may be displayed after the ATF driver is force-loaded by the operating system (this is required for DMP co-existence with the ATF driver). The messages are similar to the following:

```
WARNING: atf[ 0/ 0/ 1/ 0]:No peer found.
WARNING: atf[ 0/ 0/ 1/ 1]:No peer found.
WARNING: atf[ 0/ 0/ 1/ 2]:No peer found.
```

These messages are normal and appear only when a DG CLARiiON connected to the system is single pathed. If it is dual pathed, these messages should not appear. You can safely ignore these warnings because the co-existence of DMP and ATF is not affected by these messages.

Sun StorEdge A5x00 Array

This section pertains to the Sun StorEdge A5000, A5100, and A5200 disk arrays. These disk arrays are claimed by VxVM (DMP) under the SENA category.

When a cable connecting any of these disk arrays to the host is removed and the commands `drvconfig`, `disks`, and `vxctl enable` are executed on the host, DMP does not display the pathnames in the standard `c#t#d#s#` format. Instead, the `NONAME` string is displayed to represent the name of the removed path.

For example, assume that a host is connected to one of the A5x00 disk arrays. That disk array has two fibre cables connecting it to the host. There is one disk on the disk array. The two paths to this disk are called `/dev/[r]disk/c1t33d0s2` and `/dev/[r]disk/c2t33d0s2`. When the disk is used by VxVM, the `vxdisk list daname` command executed on this disk displays the following:

```
# vxdisk list c1t33d0s2
Device:      c1t33d0s2
devicetag:   c1t33d0
type:        auto
hostid:      gluon.veritas.com
disk:        name=rootdisk id=933470913.1042.gluon.veritas.com
group:       name=rootdg id=933470911.1025.gluon.veritas.com
```



```

info:      format=sliced, privoffset=1, pubslice=4, privslice=3
flags:    online ready autoconfig autoimport imported
pubpaths: block=/dev/vx/dmp/c1t33d0s6 char=/dev/vx/rdmp/c1t33d0s6
privpaths: block=/dev/vx/dmp/c1t33d0s7 char=/dev/vx/rdmp/c1t33d0s7
version:  2.1
iosize:   min=512 (bytes) max=2048 (blocks)
public:   slice=6 offset=1 len=35363560 disk_offset=5112
private:  slice=7 offset=1 len=4711 disk_offset=2556
update:   time=933471058 seqno=0.9
headers:  0 24 8
ssb:      actual_seqno=0.0
configs:  count=1 len=3448
logs:     count=1 len=522
Defined regions:
config   priv 000017-000247[000231]: copy=01 offset=000000 enabled
config   priv 000249-003465[003217]: copy=01 offset=000231 enabled
log      priv 003466-003987[000522]: copy=01 offset=000000 enabled
Multipathing information:
numpaths: 2
c1t33d0s2 state=enabled
c2t33d0s2 state=enabled

```

After disconnecting the cable to controller c2 and executing a `vxctl enable` command on the system, the display changes as follows:

```
# vxdisk list c1t33d0s2
```

```

Device:    c1t33d0s2
devicetag: c1t33d0
type:      auto
hostid:    gluon.veritas.com
disk:      name=rootdisk id=933470913.1042.gluon.veritas.com
group:     name=rootdg id=933470911.1025.gluon.veritas.com
info:      format=sliced, privoffset=1, pubslice=4, privslice=3
flags:    online ready autoconfig autoimport imported
pubpaths: block=/dev/vx/dmp/c1t33d0s6 char=/dev/vx/rdmp/c1t33d0s6
privpaths: block=/dev/vx/dmp/c1t33d0s7 char=/dev/vx/rdmp/c1t33d0s7
version:  2.1
iosize:   min=512 (bytes) max=2048 (blocks)
public:   slice=6 offset=1 len=35363560 disk_offset=5112
private:  slice=7 offset=1 len=4711 disk_offset=2556
update:   time=933471058 seqno=0.9
ssb:      actual_seqno=0.0
headers:  0 24 8
configs:  count=1 len=3448
logs:     count=1 len=522
Defined regions:

```



```

config priv 000017-000247[000231]: copy=01 offset=000000 enabled
config priv 000249-003465[003217]: copy=01 offset=000231 enabled
log priv 003466-003987[000522]: copy=01 offset=000000 enabled
Multipathing information:
numpaths:      2
c1t33d0s2  state=enabled
NONAME  state=disabled

```

Similarly, the output of the `vxddmpadm` command also displays the NONAME string for the names of the paths that have been removed. For example, if the controller `c1` is removed from the host, the `vxddmpadm` command displays the following:

```

# vxddmpadm getsubpaths ctlr=c1
NAME      STATE      TYPE  DMPNODENAME  ENCLR-TYPE  ENCLR-NAME
=====
NONAME    DISABLED   -     c1t36d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t49d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t38d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t37d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t53d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t34d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t51d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t54d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t52d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t33d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t50d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t32d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t48d0s2   SENA        SENA0
NONAME    DISABLED   -     c1t35d0s2   SENA        SENA0

```

Dynamic Reconfiguration

Dynamic Reconfiguration (DR) is a feature available on some high-end SUN Enterprise™ and Sun Fire™ systems. The board to be reconfigured is a system board that contains disks controlled by VxVM (in addition to CPUs, memory, and other controllers or I/O boards) that can be offlined while the system is still running. You can dynamically reconfigure your system using one of the relevant procedures described below.

Note The system board, disks, and controllers can be configured in three different ways. Each of these configurations has a different procedure and may or may not require data to be offlined.




Domain Configuration Type 1

In this configuration, each of the disk arrays connected to the system board to be detached satisfy the following criteria:

- ◆ The disk array is multipathed.
- ◆ Paths to the disk array are connected to more than one system board that constitute the *domain*. In other words, even if one board is removed, there will still be one or more access paths to the disk array, from other system boards.

An example of such a configuration is shown in the following illustration.

Domain Configuration Type 1



To reconfigure a system board in the domain, such as the System Board 8 in UE10K domain in the preceding figure, use the following steps:

1. Identify all the disk controllers that are on the system board.

To identify the full controller names of all disk controllers connected to the board being detached, use the following commands:

```
# vxddmpadm listctlr all
# drshow board_no IO
```

where *board_no* is the system board to be detached.

Note The `drshow` command is unique to the E10K System Service Processor (SSP), and therefore must be run on the SSP. `drshow` is not a Solaris Operating system command. For more information on the `drshow` command, see the Sun documentation.

For example, in Configuration 1, on the System Board 8, C1 and C2 are the disk controllers.

To obtain multipathing information, use the following command:

```
# vxdmpadm getsubpaths ctlr=ctlr-name
```

2. For each disk controller on the System Board, use the following command:

```
# vxdmpadm disable ctlr=ctlr_name
```

For example, for system board 8 shown in the figure above, use the following commands:

```
# vxdmpadm disable ctlr=C1
# vxdmpadm disable ctlr=C2
```

This command insures that I/Os are shifted over from the specified controller to an alternate controller connected to one of the other system boards (System Board 15 in the figure, above). This is a blocking command and may take some time to return to the caller depending on the number of I/Os currently active through this controller.

3. Follow the steps in the Sun documentation to dynamically reconfigure the SYSTEM and/or I/O board.
4. After the board is replaced (or attached), follow the steps in the SUN documentation to see the new drives/paths on this domain (for example, this can include the need to run utilities like `drvconfig` followed by `disks`, and so on).
5. Re-enable the Volume Manager configuration daemon using the following command:

```
# vxdctl enable
```

6. Enable all the controllers on the board being attached using the following command:

```
# vxdmpadm enable ctlr=ctlr_name
```

for each controller on the board being replaced.

For example, to attach System Board 8, run the following commands:

```
# vxdmpadm enable ctlr=C1
# vxdmpadm enable ctlr=C2
```

The commands insure that the controllers are now available for I/O to the Volume Manager.


Domain Configuration Type 2

In this case, the disk array(s) connected to the system board to be detached are not accessible via other system boards in the same domain. For example, see the following illustration.



Note This procedure may require data to be offlined during the operation if it is not on redundant VxVM volumes (example: RAID5, Mirror, and so on).

Domain Configuration Type 2



The following procedure is used for this configuration:

1. If all disks that are connected to the same disk controller exist in a single diskgroup, you should deport the diskgroup using this command:

```
# vxvg deport dgname
```
2. Run the following command on all disks:

```
# vxdisk rm daname
```
3. Follow the steps in the Sun documentation to detach the system board.
4. After the board is replaced (or attached), follow the steps in the SUN documentation to see the new drives/paths on this domain (for example, this can include the need to run utilities like `drvconfig` followed by `disks`, and so on).
5. Run the Volume Manager configuration daemon using the following command:

```
# vxvxdctl enable
```
6. When all devices are back online, import the disk group:

```
# vxvg import dgname
```
7. Restart all the volumes that were stopped as a result of the system board reconfiguration.


Domain Configuration Type 3

In this case, the system board being detached has disk arrays connected as follows:

- ◆ Some disk arrays are accessible only through the system board being reconfigured.
 - For disk arrays that are accessible through one system board, use the procedure for Domain Configuration Type 2.
- ◆ The remaining disk arrays are accessible through the other system board(s) as well.
 - For disk arrays that are accessible through multiple system boards, use the procedure for Domain Configuration Type 1.

An example configuration is shown in the following illustration:

Domain Configuration Type 3



To detach or attach System Board 8, use the following procedure:

1. For the Sun StorEdge T3 disk array which is accessible only through System Board 8, follow the procedure for Domain Configuration Type 2.
2. For the A5000 disk array which is accessible through multiple system boards, that is System Board 8 and System Board 15, follow the procedure for Domain Configuration Type 1.



Using the Sun StorEdge A3000/A3500 Array with VxVM

Note The information in this section is current at the time of this printing, but is subject to change with new versions of the Sun StorEdge A3000/A3500 Array software.

The Sun StorEdge A3000/A3500 Array subsystem with the Sun StorEdge A3000/A3500 Array RAID Manager software is compatible with VxVM 4.1. To ensure compatibility between the Sun StorEdge A3000/A3500 Array and VxVM, you must do the following:

- ◆ The Sun StorEdge A3000/A3500 Array and VxVM installation sequence must be followed exactly as documented in the section “[Installation of VxVM with the Sun StorEdge A3000/A3500 Array](#)” on page 43.
- ◆ When VxVM encapsulates a device that has a mounted file system and an entry in the `/etc/vfstab` file, VxVM updates the `/etc/vfstab` entry for that device with the VxVM device node name. VxVM is not compatible with the Sun StorEdge A3000/A3500 Array RAID Manager device node naming convention, but it is compatible with the Solaris Operating System device node naming convention.

Do not use the Sun StorEdge A3000/A3500 Array RAID Manager generated disk device nodes to encapsulate the Sun StorEdge A3000/A3500 Array devices that have been configured with file systems and have entries in the `/etc/vfstab` file, for example `/dev/RAID_module01/0s[0-7]` and `/dev/rRAID_module01/0s[0-7]`. (See the Sun StorEdge A3000/A3500 Array documentation for the device nodes.)

Instead, use the standard Solaris Operating System device nodes. For example (see the Solaris Operating System documentation for the standard device nodes):

```
/dev/rdisk/c3t4d0s0 and /dev/dsk/c3t4d0s0.
```

For example, you may have an `/etc/vfstab` entry similar to this:

```
/dev/RAID_module01/0s0 /dev/rRAID_module01/0s0 /mountpt ufs 2 \
yes -
```

In this case, use the Solaris Operating System device node that represents this device:

```
/dev/dsk/c3t4d0s0 /dev/rdisk/c3t4d0s0 /mountpt ufs 2 yes -
```

- ◆ Follow the Sun StorEdge A3000/A3500 Array controller error recovery procedures as documented by the Sun StorEdge A3000/A3500 Array subsystem documentation.

Caution Sun StorEdge A3000/A3500 Array error recovery procedures must be followed according to the Sun StorEdge A3000/A3500 Array subsystem documentation. Improper error recovery procedures can cause incompatibility with VxVM.

Installation of VxVM with the Sun StorEdge A3000/A3500 Array

The Sun StorEdge A3000/A3500 Array installation sequence must be followed as documented in the *Sun StorEdge A3000/A3500 Array Subsystem Release Notes*, the *Sun StorEdge A3000/A3500 Array System Manual*, and the *Sun StorEdge A3000/A3500 Array RAID Manager Installation Guide*.

Caution Any deviation of the installation sequence causes the StorEdge A3000/A3500 Array to be incompatible with VxVM.

VxVM should be installed only after the following steps have been completed and validated:

1. The Sun StorEdge A3000/A3500 Array subsystem is properly attached to the host computer.
2. The Sun StorEdge A3000/A3500 Array RAID Manager software is properly installed.
3. The Sun StorEdge A3000/A3500 Array Logical Unit Numbers (LUNs) are properly configured using the Sun StorEdge A3000/A3500 Array RAID Manager software.
4. The host system is rebooted using the `-r` flag (`reboot(1M)`) to rescan for attached devices upon boot, as follows:

```
# reboot -- -r
```

The `--` delimiter is used to separate the options of `reboot` from the arguments of `boot`.

5. Upon reboot, the Sun StorEdge A3000/A3500 Array RAID Manager has recognized the Sun StorEdge A3000/A3500 Array configured devices and has created the appropriate Solaris Operating system device nodes (see *VERITAS Volume Manager and Sun StorEdge A3000/A3500 Array Device Nodes*).
6. The VERITAS Volume Manager software package can now be installed.

VxVM and Sun StorEdge A3000/A3500 Array Device Nodes

After installing and configuring the Sun StorEdge A3000/A3500 Array subsystem, the host needs to be rebooted using the `-r` flag (`reboot(1M)`) to rescan for attached devices. Upon reboot, the Sun StorEdge A3000/A3500 Array scripts and Solaris Operating System startup scripts create device nodes.



After the host system reboot, the original set of device nodes should be seen (the device nodes that existed before the Sun StorEdge A3000/A3500 Array subsystem was installed), and the new device nodes for each StorEdge A3000/A3500 Array LUN. If there are more device nodes created than attached devices (including StorEdge A3000/A3500 Array LUNs), do not install VxVM until the proper device nodes exist on the host system.

Caution The Sun StorEdge A3000/A3500 Array startup scripts must properly create and configure device nodes. Improper host system device nodes are incompatible with VxVM.

Benefits in Using VxVM with the Sun StorEdge A3000/A3500 Array

This section describes some of the benefits in using the host-based VxVM software and controller-based StorEdge A3000/A3500 Array hardware. When VxVM is used in conjunction with the Sun StorEdge A3000/A3500 Array, availability, performance and manageability of the combined configuration can be improved.

Improved Manageability

The online reconfiguration capabilities of VxVM can be used to enhance the performance monitoring, tuning, and reconfiguration of the Sun StorEdge A3000/A3500 Array. These capabilities can simplify administration while increasing uptime and performance.

Performance Analysis and Online Load Balancing Between LUNs

Using the VxVM statistics-gathering capability allows administrators to analyze the I/O loads and responsiveness of volumes in the configuration and to move storage between disks. Since each LUN within the Sun StorEdge A3000/A3500 Array looks like a single disk to the host, and therefore to VxVM, it is not possible to identify the load on each spindle, or to tune the layout of the LUN itself. However, performance imbalances between arrays, or between LUNs in the same array, may be detected and alleviated. This can be done by using the VxVM performance analysis and online reconfiguration capability.

Increasing Capacity Online by LUN Concatenation

The size of a single file system or database table space is limited by the maximum size of a single LUN in a controller-based RAID subsystem. To create very large file systems or table spaces, administrators can use VxVM to combine LUNs from multiple RAID controllers into a larger “super-LUN” volume. (Refer to the chapter, *Creating Volumes*, in the *VERITAS Volume Manager Administrator’s Guide*.)

Online Data Movement Between LUNs

The backup/reconfigure/reload cycle required to change the layout of hardware RAID LUNs causes data access interruption and possible data loss. To reduce the likelihood of data loss, if extra disks are available, administrators can choose to construct the destination LUN according to the desired parameters, then copy data directly from the old location to the new one. However, even this requires the interruption of data access, since it is necessary to prevent changes to the old copy of the data after it has been copied to the new location.

If the data copying is performed by adding the new Sun StorEdge A3000/A3500 Array LUN as a VxVM mirror of the data, all writes are delivered to all mirrors, keeping all copies up-to-date at all times. Once the mirror synchronization is complete, the mirror set can be separated, and the original LUN can be removed or used for other data. The data is kept online without interrupting access throughout the operation.

Stable Backup

The method described above for data movement can also be used to provide consistent stable backup without interrupting user access. An additional mirror of data can be created, or an existing mirror can be detached from the updating applications, and that data instance can be backed-up to tape or other offline storage pool (see the *VERITAS Volume Manager Administrator's Guide*). This guarantees a backup that is internally consistent at a single point in time. Upon completion of backups, the additional mirror space can be freed, or the mirror reattached to the live data volume and resynchronized.

Dynamic Partitioning

The dynamic partitioning capability of VxVM is useful when used with the large disks presented by the system for each LUN. The Solaris Operating Systems have limits on the number of slices into which a disk can be partitioned (7 partitions plus the full-disk partition). These limits can make it impossible to divide a large LUN into the number of separate slices needed for file systems or databases.

Sun StorEdge A3000/A3500 Array Supported Configurations

VxVM supports the Sun StorEdge A3000/A3500 Array subsystem in the following configurations:

- ◆ VxVM RAID-1 (Mirroring) and Sun StorEdge A3000/A3500 Array
- ◆ VxVM RAID-0 (Striping) and Sun StorEdge A3000/A3500 Array

For more information about mirroring and striping, see the *VERITAS Volume Manager Administrator's Guide*.



Sun StorEdge A3000/A3500 Array Unsupported Configurations

The following VxVM and Sun StorEdge A3000/A3500 Array configurations are not supported:

- ◆ VxVM RAID-5 Configuration
- ◆ Sun StorEdge A3000/A3500 Array RAID-5 Configuration

Using RAID-5 at both levels in the I/O subsystem can result in poor performance for no significant gain in reliability or availability. Use of this configuration is not supported.

VxVM Hot-Relocation and Sun StorEdge A3000/A3500 Array Hot Sparing

VxVM hot-relocation allows the host system to automatically react to I/O failures on redundant (mirrored or RAID-5) VxVM objects and restore redundancy and access to these objects. Sun StorEdge A3000/A3500 Array hot sparing allows the StorEdge A3000/A3500 Array to automatically react to I/O failures internal to the StorEdge A3000/A3500 Array to restore access to StorEdge A3000/A3500 Array LUNs.

If a disk failure occurs within the Sun StorEdge A3000/A3500 Array, and the array hot sparing is configured, then the StorEdge A3000/A3500 Array provides the disk failure redundancy. If data redundancy is provided by a VxVM RAID-5 or mirror configuration, then VxVM hot-relocation can also provide disk or partial disk failure redundancy protection. The most complete level of disk redundancy is achieved with both VxVM hot-relocation and Sun StorEdge A3000/A3500 Array hot sparing enabled.

Using the Sun StorEdge T3/T3+/T4 Array with VxVM

The Sun StorEdge T3/T3+/T4 arrays can be configured in the following array modes:

- ◆ RW (implicit failover or auto-trespass) mode.
LUN failover will be triggered on any read or write command with the exception of a read of LBA 0. To enable the implicit LUN failover, you should use the `sys mp_support=rw` command on the Sun StorEdge T3/T3+/T4 arrays.
- ◆ MPxIO (explicit failover or nonauto-trespass) mode.
LUN failover will be triggered with an array specific command. To enable the explicit LUN failover, you should use the `sys mp_support=mpxio` on the Sun StorEdge T3/T3+/T4 arrays.

In this mode, the Sun StorEdge Traffic Manager (SSTM) driver or DMP can be used to handle multipathing.

If you want to enable DMP to handle multipathing and your system is running Solaris 8 or Solaris 9, the `/kernel/drv/scsi_vhci.conf` file must have the MPxIO mode enabled (`mpxio-disable="yes"`).

If you want to enable DMP to handle multipathing and your system is running Solaris 10, the `/kernel/drv/fp.conf` file must have the MPxIO mode enabled (`mpxio-disable="yes"`).

Note All hosts in the clustered environment must use the same configuration set-up for the T3/T3+/T4 array.

Working with VxVM Disks on a Sun StorEdge T3 Array

Using a T3+ as a Boot Device

To use a T3 as a boot device, the T3 needs to satisfy some requirements:

1. All T3 firmware should be at level version 2.1 or higher. To confirm that the T3 firmware is at the necessary level, log on to the T3+ and enter the command:

```
# t3plus:/:<1>ver
```

and you should see output similar to the following:

```
T3B Release 2.01.00 2002/03/22
```

2. The `volslice` parameter should be in a disabled state. Log on to the device and enter the command:

```
# t3plus:/:<1>sys list
```

you will see output, similar to the following, that includes a reference to the `enable_volslice`:

```
enable_volslice : on
```

If `volslice` is set to on, you cannot boot from the T3+ device.

The T3+ device should be directly connected to the machine, and no switch should be involved. If you are connecting through a switch, then you need to refer to the Solaris Operating system and T3+ documentation for booting requirements.



T3B firmware upgrade Considerations

For the rolling upgrade to work, you need to break the mirrors and remove the disks from the disk group. The upgrade can then be carried out on any array that is not in VM. The disks are then added back to the disk group and the mirrors re-attached. This requires that the mirrors are resynchronized, as in any mirror detach plex.

The following scenario shows how to upgrade T3B firmware from 2.0 to 2.1.

Note It is required that a volume has one plex on each T3 unit. If the volume is not mirrored across T3 units, please contact your hardware vendor.

1. Identify the volume which is on the T3B which you are going to upgrade the Firmware. For example:

```
# vxprint -ht
```

```
dg new-dg          default      default  129000
1016220570.1347.csha5.veritas.com
dm new-dg01        c0t1d0s2    sliced   4711     35363560 -
dm new-dg02        c2t1d0s2    sliced   6143     212146176 -
v mirror          -           ENABLED  ACTIVE   20480    SELECT   -

fsgen

pl mirror-01      mirror ENABLED  ACTIVE   23560    CONCAT   - RW
sd new-dg01-02    mirror-01 new-dg01 23560    23560    0 c0t1d0 ENA
pl mirror-02      mirror ENABLED  ACTIVE   24576    CONCAT   - RW
sd new-dg02-03    mirror-02 new-dg02 30720    24576    0 c2t1d0 ENA
```

2. Destroy one set of plexes from the above volume (destroy the plex which is on the T3B disk to be upgraded):

```
# vxplex -g new-dg -o rm dis mirror-02
```

```
# vxprint -ht
```

```
dg new-dg          default      default  129000
1016220570.1347.csha5.veritas.com

dm new-dg01        c0t1d0s2    sliced   4711     35363560 -
dm new-dg02        c2t1d0s2    sliced   6143     212146176 -

v mirror -         ENABLED  ACTIVE   20480    SELECT   -

fsgen

pl mirror-01      mirror          ENABLED  ACTIVE   23560    CONCAT   - RW
sd new-dg01-02    mirror-01      new-dg01 23560    23560    0 c0t1d0 ENA
```



3. Remove the disk[s] from VM control:

```
# vxdg -g new-dg rmdisk new-dg02
# vxprint -ht

dg new-dg          default          default  129000
1016220570.1347.csha5.veritas.com

dm new-dg01        c0t1d0s2         sliced   4711      35363560 -
v mirror          -                ENABLED  ACTIVE    20480     SELECT    -
fsgen

pl mirror-01       mirror           ENABLED  ACTIVE    23560     CONCAT - RW
sd new-dg01-02    mirror-01        new-dg01 23560     23560     0 c0t1d0 ENA
```

4. There should now be no disks from the T3B array in any disk group under VM control. If any disks from the T3B array are in any disk group, return to Steps 1 through 3 to remove devices from VM control.**5. Follow the upgrade procedures to upgrade T3B firmware from 2.0 to 2.1, as specified in the upgrade documentation of the hardware.****6. Run `vxdctl enable` to make device available back to VM:**

```
# vxdctl enable
```

7. Add the disk to original disk group and use original name:

```
# vxdiskadd c2t1d0

or

# vxdg -g new-dg adddisk new-dg02=c2t1d0
# vxprint -ht

dg new-dg          default          default  129000
1016220570.1347.csha5.veritas.com

dm new-dg01        c0t1d0s2         sliced   4711      35363560 -
dm new-dg02        c2t1d0s2         sliced   6143      212146176 -
v mirror          -                ENABLED  ACTIVE    20480     SELECT    -
fsgen

pl mirror-01       mirror           ENABLED  ACTIVE    23560     CONCAT - RW
sd new-dg01-02    mirror-01        new-dg01 23560     23560     0 c0t1d0 ENA
```



8. Re-create the mirror:

```
# vxassist -g new-dg mirror mirror new-dg02
# vxprint -ht

dg new-dg          default          default  129000
1016220570.1347.csha5.veritas.com

dm new-dg01        c0t1d0s2         sliced   4711    35363560 -
dm new-dg02        c2t1d0s2         sliced   6143    212146176 -

v mirror           -                ENABLED  ACTIVE  20480    SELECT   -

fsgen

pl mirror-01       mirror           ENABLED  ACTIVE  23560    CONCAT - RW
sd new-dg01-02    mirror-01        new-dg01 23560    23560    0 c0t1d0 ENA
pl mirror-02       mirror           ENABLED  TEMPRMSD 24576    CONCAT - RO
sd new-dg02-01    mirror-02        new-dg02 0         24576    0 c2t1d0 ENA
```

9. Wait until the mirrors are in sync for the volume. This can be determined by examining the plex's State field in vxprint output. Once the mirrors are in sync, you will see the following in the vxprint output:

```
# vxprint -ht

dg new-dg          default          default  129000
1016220570.1347.csha5.veritas.com

dm new-dg01        c0t1d0s2         sliced   4711    35363560 -
dm new-dg02        c2t1d0s2         sliced   6143    212146176 -

v mirror           -                ENABLED  ACTIVE  20480    SELECT   -

fsgen

pl mirror-01       mirror           ENABLED  ACTIVE  23560    CONCAT - RW
sd new-dg01-02    mirror-01        new-dg01 23560    23560    0 c0t1d0 ENA

fsgen

pl mirror-02       mirror           ENABLED  ACTIVE  24576    CONCAT - RW
sd new-dg02-01    mirror-02        new-dg02 0         24576    0 c2t1d0 ENA
```

Follow the above steps to perform an upgrade on another T3 unit.



Index

A

- array support libraries
 - adding 15
 - installing 14
 - older releases 17
 - removing 16
 - supported 3
 - upgrading 17
- arrays 29
 - HDS 5700E 29
 - HDS 5800E/9200 30
 - HDS 7700E 30
 - HP surestore E disk array XP25 22
 - HP Surestore E disk array XP512 24
 - IBM Enterprise Storage Server (ESS) 26
 - Sun StorEdge 6120/6320 17
 - Sun StorEdge A3000/A3500 38, 42
 - benefits 44
 - hot-relocation 46
 - installation of VxVM 43
 - supported configurations 45
 - unsupported configurations 46
 - Sun StorEdge A5x00 35
 - supported for boot mirroring 5

D

- disks
 - partitioning on Sun StorEdge 6120/6320 arrays 17
- DMP
 - administering 10
 - coexistence with clariion ATF driver 35
 - JBOD disk array support 31
 - supported disk arrays 2
 - using the AP driver 32
 - using the Clariion ATF driver 34
 - vxdmpadm 10
 - working with EMC symmetrix disk array 21

- domain configuration 38
- dynamic reconfiguration 37

E

- extended copy service
 - enabling a disk 18
 - introduction 18

H

- HDS 5700E 29
- HDS 5800E 30
- HDS 7700E 30
- HDS 9200 30
- Hitachi arrays 29
- host bus adapters
 - JNI 20
- HP SureStore E disk array XP256 22, 24
- HP SureStore E disk array XP512 24

I

- IBM Enterprise storage server 26
- IBM subsystem device driver (SDD) 29

L

- LUN expansion 12

P

- prestoserve
 - using with VxVM 5

S

- Sun StorEdge 6120/6320 17
- Sun StorEdge A3000/A3500 42
 - array device nodes 43
 - benefits 44
 - hot relocation 46
 - installation of VxVM 43
 - supported configurations 45
 - unsupported configurations 46
- Sun StorEdge A5x00 35



T

Technical assistance ix

