

VERITAS Volume Manager™ 3.2

Hardware Notes

Solaris

August 2001
30-000397-011


VERITAS

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.

Copyright

Copyright © 2000-2001 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the U.S. and other countries. The VERITAS logo and VERITAS Volume Manager are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA/Ireland, April 2001.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
Phone 650-527-8000
Fax 650-527-2901
www.veritas.com



Contents

Preface	vii
Introduction	vii
Audience and Scope	vii
Organization	vii
Related Documents	viii
Conventions	ix
Getting Help	x
Downloading and Running VRTSexplorer	x
Hardware Notes	1
Introduction	1
New Hardware Support in VxVM 3.2	1
Using Prestoserve With VxVM	1
The DMP Command Utility (vxdkmpadm)	4
VxVM Disk Array Configuration Updates	6
Disk Arrays Supported for DMP	6
DMP and EMC Symmetrix Disk Array Support	7
HP SureStore™ E Disk Array XP256	8
HP SureStore™ E Disk Array XP512	10
IBM Enterprise Storage Server (ESS)	12
DMP and Hitachi Data Systems 5700E SCSI Disk Array Subsystem Support ..	15
DMP and Hitachi Data Systems 5800E / 7700E / 9910 / 9960 / 9200 Fiber Disk Array Subsystem Support	15
DMP and JBOD Disk Array Support	16



DMP Co-existence With Drivers	17
Using the AP Driver With the VERITAS Volume Manager	18
Using the ATF Driver With VxVM	20
DMP and ATF Driver Co-existence Procedure	21
Sun StorEdge A5x00 Array	21
Minimum Revision Levels for Solaris Patches	24
Dynamic Reconfiguration	24
Domain Configuration Type 1	25
Domain Configuration Type 2	27
Domain Configuration Type 3	28
Using the Sun StorEdge A3000/A3500 Array With VxVM	29
Installation of VxVM With the Sun StorEdge A3000/A3500 Array	30
VxVM and Sun StorEdge A3000/A3500 Array Device Nodes	31
Benefits in Using VxVM With the Sun StorEdge A3000/A3500 Array	31
Improved Manageability	31
Performance Analysis and Online Load Balancing Between LUNs	31
Increasing Capacity Online by LUN Concatenation	32
Online Data Movement Between LUNs	32
Stable Backup	32
Dynamic Partitioning	32
Sun StorEdge A3000/A3500 Array Supported Configurations	33
Sun StorEdge A3000/A3500 Array Unsupported Configurations	33
VxVM Hot-Relocation and Sun StorEdge A3000/A3500 Array Hot Sparring	33
Using the Sun StorEdge T3 Array With VxVM	34
Sun StorEdge T3 Support	34
Sun StorEdge T3 Support In A Single Host Environment	34
Sun StorEdge T3 Support In A Multi-host Environment	34
Working with VxVM Disks on a Sun StorEdge T3	34
Initializing a Disk on a T3 or T3+ Array in c##t##d# Mode	35
Removing or Uninitializing a Disk on a T3 or T3+ Array in c##t##d# Mode	35



Replacing a Disk on a T3 or T3+ Array in c#t#d# Mode	35
--	----





Preface

Introduction

The *VERITAS Volume Manager™ Hardware Notes* provides hardware information on VERITAS Volume Manager (VxVM) Release 3.2.

Audience and Scope

This guide is intended for system administrators responsible for installing, configuring, and maintaining systems under the control of VxVM. The purpose of this guide is to provide the system administrator with information on hardware requirements for VERITAS Volume Manager Release 3.2.

This guide assumes that the user has a:

- ◆ working knowledge of the UNIX operating system
- ◆ basic understanding of UNIX system administration
- ◆ basic understanding of volume management

Organization

This guide is organized as follows:

- ◆ [Using Prestoserve With VxVM](#)
- ◆ [VxVM Disk Array Configuration Updates](#)
- ◆ [Disk Arrays Supported for DMP](#)



Related Documents

The following documents provide information related to VxVM:

- ◆ *VERITAS Volume Manager Installation Guide*
- ◆ *VERITAS Volume Manager Troubleshooting Guide*
- ◆ *VERITAS Volume Manager Release Notes*
- ◆ *VERITAS Volume Manager Administrator's Guide*
- ◆ *VERITAS Volume Manager Storage Administrator Administrator's Guide*
- ◆ VERITAS Volume Manager manual pages



Conventions

The following table describes the typographic conventions used in this guide.

Typeface	Usage	Examples
monospace	Computer output, file contents, files, directories, software elements such as command options, function names, and parameters	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls(1)</code> manual page for more information.
<i>italic</i>	New terms, book titles, emphasis, variables to be replaced by a name or value	See the <i>User's Guide</i> for details. The variable <code>ncsize</code> determines the value of...
monospace (bold)	User input; the “#” symbol indicates a command prompt	<code># mount -F vxfs /h/filesys</code>
<i>monospace (bold and italic)</i>	Variables to be replaced by a name or value in user input	<code># mount -F <i>fstype mount_point</i></code>

Symbol	Usage	Examples
%	C shell prompt	
\$	Bourne/Korn/Bash shell prompt	
#	Superuser prompt (all shells)	
\	Continued input on the following line	<code># mount -F vxfs \ /h/filesys</code>
[]	In a command synopsis, brackets indicates an optional argument	<code>ls [-a]</code>
	In a command synopsis, a vertical bar separates mutually exclusive arguments	<code>mount [suid nosuid]</code>



Getting Help

If you have any comments or problems with the VERITAS products, contact the VERITAS Technical Support:

- ◆ U.S. and Canadian Customers: 1-800-342-0652
- ◆ International Customers: +1 (650) 527-8555
- ◆ Email: support@veritas.com

For license information (U.S. and Canadian Customers):

- ◆ Phone: 1-925-931-2464
- ◆ Email: license@veritas.com
- ◆ Fax: 1-925-931-2487

For software updates:

- ◆ Email: swupdate@veritas.com

For additional technical support information, such as TechNotes, product alerts, and hardware compatibility lists, visit the VERITAS Technical Support Web site at:

- ◆ <http://support.veritas.com> (U.S. and Canada Customers)

For additional information about VERITAS and VERITAS products, visit the Web site at:

- ◆ <http://www.veritas.com>

Downloading and Running VRTSexplorer

If you have access to the Internet, you can use the VRTSexplorer program to assist Technical Support in diagnosing the cause of your problem as follows:

1. Use a web browser or the `ftp` program to download the VRTSexplorer program at the following URL:

```
ftp://ftp.veritas.com/pub/support/vxexplore.tar.Z
```

Save the file to a temporary directory such as `/tmp` as shown in these instructions. If you download the file to a different directory, substitute its pathname for `/tmp` throughout.

2. Log in as `root` on the affected system, and use the following commands to extract the contents of the downloaded file to the directory `/tmp/VRTSexplorer`:

```
# cd /tmp
# zcat vxexplore.tar.Z | tar xvf -
```

3. Run the `VRTSexplorer` program located in the `VRTSexplorer` directory by entering the following command:

```
# /tmp/VRTSexplorer/VRTSexplorer
```

4. When `VRTSexplorer` prompts you for a destination directory for the information that it collects, press Return to accept the default directory `/tmp`, or enter a pathname of your own choice. `VRTSexplorer` writes the results of its investigations to a compressed tar file named `VRTSexplorer.case_number.tar.Z` in the specified directory.
5. Use the file upload facility of your web browser or the `ftp` program to transfer the file output by `VRTSexplorer` to the VERITAS Customer Support anonymous FTP site:

```
ftp://ftp.veritas.com/incoming
```

6. Call VERITAS Customer Support on 1-800-342-0652, inform them that you have run `VRTSexplorer` and tell them the name of the file that you transferred to the FTP site.

Alternatively, if you have already been assigned a call ID number by Customer Support, e-mail `support@veritas.com` including your case ID number in the subject line.

For more information about the `VRTSexplorer` program, consult the `README` file located in the `VRTSexplorer` directory.





Hardware Notes

Introduction

This document provides hardware support information for VERITAS Volume Manager™ (VxVM) Release 3.2.

This release of VxVM has been tested on:

- Solaris 2.6
- Solaris 7 (32- and 64-bit)
- Solaris 8 (32- and 64-bit).

New Hardware Support in VxVM 3.2

VxVM 3.2 now supports the following hardware:

- ◆ SunFire 3800, 4800, and 6800 systems.
- ◆ NEC arrays: DS1200, DS1200F, DS3000SL, DS3000SM, DS3001, DS3002, DS1000, DS1000F, DS1100, DS1100F, DS3011, DS1230, DS450, DS450F, iStorage 1000, iStorage 2000, and iStorage 4000.
- ◆ HDS 9910/9960/9200 disk array.
- ◆ HP SureStore E Disk Array XP48.

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 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. 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
```

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. Anytime a volume device number is temporarily remapped, the following console message displays:

```
vxvm:vxconfigd: WARNING: 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 `vxdg remminor` command.

One method for ensuring that Prestoserve never encounters two volumes with conflicting device numbers is to use Prestoserve only on volumes in the `rootdg` disk group. The `rootdg` disk group is always imported first, so it cannot conflict with a previously-imported disk group.



The DMP Command Utility (vxdumpadm)

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

You can use the `vxdumpadm` utility command `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:

```
# vxdumpadm listctlr all
```

to display the following example output:

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c0	OTHER	ENABLED	others0
c1	SEAGATE	ENABLED	seagate0
c2	SEAGATE	ENABLED	seagate0

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 A5x00 (SEAGATE) disk array

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

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

```
# vxdumpadm getsubpaths ctlr=c1
```

to display the following example output:

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c1t0d0s2	ENABLED	-	c2t0d0s2	SEAGATE	seagate0
c1t1d0s2	ENABLED	-	c2t1d0s2	SEAGATE	seagate0
c1t2d0s2	ENABLED	-	c2t2d0s2	SEAGATE	seagate0
c1t3d0s2	ENABLED	-	c2t3d0s2	SEAGATE	seagate0
c1t4d0s2	ENABLED	-	c2t4d0s2	SEAGATE	seagate0
c1t5d0s2	ENABLED	-	c2t5d0s2	SEAGATE	seagate0
c1t6d0s2	ENABLED	-	c2t6d0s2	SEAGATE	seagate0

c1t16d0s2	ENABLED	-	c2t16d0s2	SEAGATE	seagate0
c1t17d0s2	ENABLED	-	c2t17d0s2	SEAGATE	seagate0
c1t18d0s2	ENABLED	-	c2t18d0s2	SEAGATE	seagate0
c1t19d0s2	ENABLED	-	c2t19d0s2	SEAGATE	seagate0
c1t20d0s2	ENABLED	-	c2t20d0s2	SEAGATE	seagate0
c1t21d0s2	ENABLED	-	c2t21d0s2	SEAGATE	seagate0
c1t22d0s2	ENABLED	-	c2t22d0s2	SEAGATE	seagate0

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 `c1t0d0s2` 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 following example output:

NAME	STATE	PATH-TYPE	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME
c2t0d0s2	DISABLED	-	c2	SEAGATE	seagate0
c1t0d0s2	ENABLED	-	c1	SEAGATE	seagate0

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 following example output:

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
c2t0d0s2	ENABLED	SEAGATE	2	2	0	seagate0

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



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. These topics are described in this section.

- ◆ [Disk Arrays Supported for DMP](#)
- ◆ [DMP Co-existence With Drivers](#)
- ◆ [Sun StorEdge A5x00 Array](#)
- ◆ [Minimum Revision Levels for Solaris Patches](#)
- ◆ [Dynamic Reconfiguration](#)
- ◆ [Using the Sun StorEdge T3 Array With VxVM](#)

Disk Arrays Supported for DMP

Note Arrays supported by the DMP feature are a subset of the arrays supported by VxVM. The following arrays can have multiple paths to the same disk. Other arrays not listed here can only have a single path to the physical disk.

The DMP feature of VxVM supports the following arrays:

- ◆ EMC Symmetrix™
- ◆ HP SureStore™ E Disk Array XP256/XP512/ XP48
- ◆ IBM Enterprise Storage Servers™ (ESS)
- ◆ Hitachi Data Systems™ 5700E Disk Array Subsystem™
- ◆ Hitachi Data Systems 5800E/7700E Disk Array Subsystem™
- ◆ Hitachi Data Systems 9910, 9960 and 9200 Disk Array Subsystem™
- ◆ Sun StorEdge A5x00 Array™
- ◆ Sun StorEdge T3 Array™
- ◆ JBOD (Just a Bunch of Disks)
- ◆ NEC arrays: DS1200, DS1200F, DS3000SL, DS3000SM, DS3001, DS3002, DS1000, DS1000F, DS1100, DS1100F, DS3011, DS1230, DS450, DS450F, iStorage 1000, iStorage 2000, 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 3.2, you must assign a unique system name for each disk array connected to a machine. The *RAID-5 Users Manual* at the web site www.storage.com describes how to set a system name for Storage Computer disk arrays.
- ◆ ECCS™ Synchronix™ 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* line 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 and setting this value reduces the time considerably.

See 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:

```
# vxddmpadm listctlr all
```

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c0	OTHER	ENABLED	others0
c5	XP256	ENABLED	xp256-0
c6	XP256	DISABLED	xp256-0

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

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c5t11d1s2	ENABLED	-	c5t11d1s2	XP256	xp256-0
c5t11d2s2	ENABLED	-	c5t11d2s2	XP256	xp256-0
c5t11d3s2	ENABLED	-	c5t11d3s2	XP256	xp256-0
c5t11d4s2	ENABLED	-	c5t11d4s2	XP256	xp256-0
c5t11d5s2	ENABLED	-	c5t11d5s2	XP256	xp256-0
c5t11d6s2	ENABLED	-	c5t11d6s2	XP256	xp256-0
c5t11d7s2	ENABLED	-	c5t11d7s2	XP256	xp256-0
c5t11d8s2	ENABLED	-	c5t11d8s2	XP256	xp256-0
c5t11d9s2	ENABLED	-	c5t11d9s2	XP256	xp256-0
c5t11d10s2	ENABLED	-	c5t11d10s2	XP256	xp256-0
c5t11d11s2	ENABLED	-	c5t11d11s2	XP256	xp256-0

```
c5t11d12s2 ENABLED - c5t11d12s2 XP256 xp256-0
```

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	sliced	-	-	error
c5t11d1s2	sliced	disk04	rootdg	online
c5t11d2s2	sliced	disk05	rootdg	online
c5t11d3s2	sliced	disk06	rootdg	online
c5t11d4s2	sliced	disk07	rootdg	online
c5t11d5s2	sliced	disk08	rootdg	online
c5t11d6s2	sliced	-	-	error
c5t11d7s2	sliced	-	-	error
c5t11d8s2	sliced	-	-	error
c5t11d9s2	sliced	-	-	error
c5t11d10s2	sliced	disk01	rootdg	online
c5t11d11s2	sliced	disk02	rootdg	online
c5t11d12s2	sliced	disk03	rootdg	online

```
# vxdisk list disk03
```

```
Device:      c5t11d12s2
devicetag:  c5t11d12
type:       sliced
hostid:     romance
disk:       name=disk03 id=961801990.1071.romance
group:      name=rootdg id=961801968.1025.romance
flags:      online ready private autoconfig autoimport imported
pubpaths:   block=/dev/vx/dmp/c5t11d12s4
             char=/dev/vx/rdmp/c5t11d12s4
privpaths:  block=/dev/vx/dmp/c5t11d12s3
             char=/dev/vx/rdmp/c5t11d12s3
version:    2.1
iosize:     min=512 (bytes) max=256 (blocks)
public:     slice=4 offset=0 len=14417280
private:    slice=3 offset=1 len=2879
update:     time=963519037 seqno=0.240
headers:    0 248
configs:    count=1 len=2104
logs:       count=1 len=318
Defined regions:
  config priv 000017-000247[000231]: copy=01 offset=000000
  disabled
  config priv 000249-002121[001873]: copy=01 offset=000231
  disabled
  log priv 002122-002439[000318]: copy=01 offset=000000
  disabled
Multipathing information:
```



```

numpaths:      2
c5t11d12s2     state=enabled
c6t8d12s2     state=disabled

```

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.

Output examples include:

```
# vxdmpadm listctlr all
```

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c0	SEAGATE	ENABLED	seagate0
c2	SEAGATE	ENABLED	seagate0
c5	XP512	ENABLED	xp512-0
c6	XP512	ENABLED	xp512-0

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

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c5t0d0s2	ENABLED	-	c5t0d0s2	XP512	xp512-0
c5t0d1s2	ENABLED	-	c5t0d1s2	XP512	xp512-0
c5t0d2s2	ENABLED	-	c5t0d2s2	XP512	xp512-0
c5t0d3s2	ENABLED	-	c5t0d3s2	XP512	xp512-0
c5t0d4s2	ENABLED	-	c5t0d4s2	XP512	xp512-0
c5t0d5s2	ENABLED	-	c5t0d5s2	XP512	xp512-0
c5t0d6s2	ENABLED	-	c5t0d6s2	XP512	xp512-0
c5t0d7s2	ENABLED	-	c5t0d7s2	XP512	xp512-0
c5t0d9s2	ENABLED	-	c5t0d9s2	XP512	xp512-0
c5t0d10s2	ENABLED	-	c5t0d10s2	XP512	xp512-0
c5t0d11s2	ENABLED	-	c5t0d11s2	XP512	xp512-0
c5t0d12s2	ENABLED	-	c5t0d12s2	XP512	xp512-0
c5t0d13s2	ENABLED	-	c5t0d13s2	XP512	xp512-0
c5t0d14s2	ENABLED	-	c5t0d14s2	XP512	xp512-0
c5t0d15s2	ENABLED	-	c5t0d15s2	XP512	xp512-0

```
# vxdmpadm getsubpaths ctlr=c6
```

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c6t0d0s2	ENABLED	-	c5t0d0s2	XP512	xp512-0



```

c6t0d1s2      ENABLED      -           c5t0d1s2      XP512         xp512-0
c6t0d2s2      ENABLED      -           c5t0d2s2      XP512         xp512-0
c6t0d3s2      ENABLED      -           c5t0d3s2      XP512         xp512-0
c6t0d4s2      ENABLED      -           c5t0d4s2      XP512         xp512-0
c6t0d5s2      ENABLED      -           c5t0d5s2      XP512         xp512-0
c6t0d6s2      ENABLED      -           c5t0d6s2      XP512         xp512-0
c6t0d7s2      ENABLED      -           c5t0d7s2      XP512         xp512-0
c6t0d9s2      ENABLED      -           c5t0d9s2      XP512         xp512-0
c6t0d10s2     ENABLED      -           c5t0d10s2     XP512         xp512-0
c6t0d11s2     ENABLED      -           c5t0d11s2     XP512         xp512-0
c6t0d12s2     ENABLED      -           c5t0d12s2     XP512         xp512-0
c6t0d13s2     ENABLED      -           c5t0d13s2     XP512         xp512-0
c6t0d14s2     ENABLED      -           c5t0d14s2     XP512         xp512-0
c6t0d15s2     ENABLED      -           c5t0d15s2     XP512         xp512-0

```

vxdisk list

DEVICE	TYPE	DISK	GROUP	STATUS
c0t4d0s2	sliced	-	-	error
c0t5d0s2	sliced	-	-	error
c0t6d0s2	sliced	-	-	error
c0t7d0s2	sliced	-	-	error
c2t0d0s2	sliced	-	-	error
c2t1d0s2	sliced	-	-	error
c2t2d0s2	sliced	-	-	error
c2t3d0s2	sliced	-	-	error
c5t0d0s2	sliced	disk01	rootdg	online
c5t0d1s2	sliced	disk02	rootdg	online
c5t0d2s2	sliced	disk03	rootdg	online
c5t0d3s2	sliced	disk04	rootdg	online
c5t0d4s2	sliced	disk05	rootdg	online
c5t0d5s2	sliced	disk06	rootdg	online
c5t0d6s2	sliced	disk07	rootdg	online
c5t0d7s2	sliced	disk08	rootdg	online
c5t0d9s2	sliced	-	-	error
c5t0d10s2	sliced	-	-	error
c5t0d11s2	sliced	-	-	error
c5t0d12s2	sliced	-	-	error
c5t0d13s2	sliced	-	-	error
c5t0d14s2	sliced	-	-	error
c5t0d15s2	sliced	-	-	online

vxdisk list c5t0d1s2

```

Device:      c5t0d1s2
devicetag:  c5t0d1
type:       sliced

```



```

hostid:      kirin
disk:        name=disk02 id=976924352.1097.kirin
group:       name=rootdg id=976924305.1025.kirin
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
private:     slice=3 offset=1 len=2879
update:      time=976952456 seqno=0.88
headers:     0 248
configs:     count=1 len=2104
logs:        count=1 len=318
Defined regions:
  config  priv 000017-000247[000231]: copy=01 offset=000000 enabled
  config  priv 000249-002121[001873]: copy=01 offset=000231 enabled
  log     priv 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 parameters for JNI HBA are:

- ◆ `set sd:sd_max_throttle = 24`
- ◆ `set sd:sd_io_time=0xf0`
- ◆ `set sd:sd_retry_count=10`

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          ENCLR-NAME

```




```

=====
c0                OTHER                ENABLED    others0
c3                IBM-SHARK                ENABLED    shark0
c4                IBM-SHARK                DISABLED   shark0
c0                SEAGATE                 ENABLED    seagate0

# vxddmpadm getsubpaths ctlr=c4

NAME              STATE          PATH-TYPE    DMPNODENAME   ENCLR-TYPE    ENCLR-NAME
=====
c4t1d2s2         DISABLED      -            c3t0d2s2      IBM-SHARK     shark0
c4t1d3s2         DISABLED      -            c4t1d3s2      IBM-SHARK     shark0
c4t1d4s2         DISABLED      -            c3t0d4s2      IBM-SHARK     shark0
c4t1d5s2         DISABLED      -            c3t0d5s2      IBM-SHARK     shark0
c4t1d6s2         DISABLED      -            c3t0d6s2      IBM-SHARK     shark0
c4t1d7s2         DISABLED      -            c3t0d7s2      IBM-SHARK     shark0
c4t1d8s2         DISABLED      -            c3t0d8s2      IBM-SHARK     shark0
c4t1d9s2         DISABLED      -            c3t0d9s2      IBM-SHARK     shark0
c4t1d10s2        DISABLED      -            c3t0d10s2     IBM-SHARK     shark0
c4t1d11s2        DISABLED      -            c3t0d11s2     IBM-SHARK     shark0
c4t1d12s2        DISABLED      -            c3t0d12s2     IBM-SHARK     shark0
c4t1d13s2        DISABLED      -            c3t0d13s2     IBM-SHARK     shark0
c4t1d14s2        DISABLED      -            c3t0d14s2     IBM-SHARK     shark0

# vxddisk list

DEVICE           TYPE          DISK          GROUP          STATUS
c0t0d0s2         sliced        -             -              error
c0t1d0s2         sliced        -             -              error
c3t0d1s2         sliced        disk06        rootdg         online
c3t0d2s2         sliced        disk07        rootdg         online
c3t0d4s2         sliced        disk09        rootdg         online
c3t0d5s2         sliced        disk10        rootdg         online
c3t0d6s2         sliced        disk11        rootdg         online
c3t0d7s2         sliced        disk14        rootdg         online
c3t0d8s2         sliced        disk12        rootdg         online
c3t0d9s2         sliced        disk13        rootdg         online
c3t0d10s2        sliced        -             -              error
c3t0d11s2        sliced        -             -              error
c3t0d12s2        sliced        -             -              error
c3t0d13s2        sliced        -             -              error
c3t0d14s2        sliced        -             -              error
c3t1d10s6        simple        -             -              error
c3t1d11s6        simple        -             -              error
c4t1d3s2         sliced        disk08        rootdg         online

# vxddisk list disk08

Device:          c4t1d3s2

```



```

devicetag: c4t1d3
type:      sliced
hostid:    banjo
disk:      name=disk08 id=963015532.1097.banjo
group:     name=rootdg id=963015492.1025.banjo
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
private:   slice=3 offset=1 len=3839
update:    time=963517507 seqno=0.213
headers:   0 248
configs:   count=1 len=2808
logs:      count=1 len=425
Defined regions:
  config  priv 000017-000247[000231]: copy=01 offset=000000 enabled
  config  priv 000249-002825[002577]: copy=01 offset=000231 enabled
  log     priv 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	others0
c3	IBM-SHARK	ENABLED	shark0
c4	IBM-SHARK	DISABLED	shark0
c0	SEAGATE	ENABLED	seagate0

vxddmpadm getsubpaths ctlr=c3

NAME	STATE	PATH-TYPE	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME
c3t0d1s2	ENABLED	-	c3t0d1s2	IBM-SHARK	shark0
c3t0d2s2	ENABLED	-	c3t0d2s2	IBM-SHARK	shark0
c3t0d4s2	ENABLED	-	c3t0d4s2	IBM-SHARK	shark0
c3t0d5s2	ENABLED	-	c3t0d5s2	IBM-SHARK	shark0
c3t0d6s2	ENABLED	-	c3t0d6s2	IBM-SHARK	shark0
c3t0d7s2	ENABLED	-	c3t0d7s2	IBM-SHARK	shark0
c3t0d8s2	ENABLED	-	c3t0d8s2	IBM-SHARK	shark0
c3t0d9s2	ENABLED	-	c3t0d9s2	IBM-SHARK	shark0
c3t0d10s2	ENABLED	-	c3t0d10s2	IBM-SHARK	shark0
c3t0d11s2	ENABLED	-	c3t0d11s2	IBM-SHARK	shark0



c3t0d12s2	ENABLED	-	c3t0d12s2	IBM-SHARK	shark0
c3t0d13s2	ENABLED	-	c3t0d13s2	IBM-SHARK	shark0
c3t0d14s2	ENABLED	-	c3t0d14s2	IBM-SHARK	shark0
c3t0d3s2	ENABLED	-	c4t1d3s2	IBM-SHARK	shark0

DMP and Hitachi Data Systems 5700E SCSI Disk Array Subsystem Support

DMP supports the Hitachi Data Systems 5700E Disk Array Subsystem 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
DF400 5800	x457/H	a03, b13

Contact your Hitachi Data Systems Technical Support Representative for instructions.

DMP and Hitachi Data Systems 5800E / 7700E / 9910 / 9960 / 9200 Fiber Disk Array Subsystem Support

DMP supports the Hitachi Data Systems 5800E/7700E Disk Array Subsystems in *Active/Active* mode only. The 5800E/7700E are treated as a *jbod* by DMP when the string **HITACHI** is added to the `/kernel/drv/vxdmp.conf` file in the `dmp_jbods[]` array. This entry is present in the `/kernel/drv/vxdmp.conf` file by default.



If more than 500 volumes are to be created in a single disk group, the private region for each disk must be set to at least 1 megabyte. This can be accomplished by initializing each disk with the `vxdisksetup` command, as shown here:

```
# vxdisksetup -i c#t#d# privlen=size
```

The value of *size* should be at least 1m.

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 unique serial numbers on disks.

- ◆ The JBOD disk arrays that require multipathing functionality from DMP should be of the Active/Active type. This means that all paths to all LUNs on the JBOD disk array should be simultaneously accessible through all paths, without causing any performance degradation.

Once it is known that disks connected to a JBOD disk array satisfy the above conditions, the following four steps must be followed to get these disks under the JBOD category of DMP:

1. Go to the `/kernel/drv` directory
2. Edit the file `vxddmp.conf` in that directory
3. You will see the following default entry in the `vxddmp.conf`:

```
dmp_jbod= "HITACHI "
```

For example, assume that there is a JBOD connected to a machine that contains DG Clariion disks. You need to know the exact Vendor ID string returned by these disks to configure DMP. This is to ensure that the required disks are claimed under the JBOD category.

The `format` command can be used to retrieve the Vendor ID string for a disk (for more details about the `format` command, check the `format(1M)` manual page on the system). The Vendor ID string is contained in bytes 8—15 of standard SCSI INQUIRY data returned by the disk. The `inquiry` option of the `format` command (when executed on a disk) returns the Vendor ID associated with a disk under the `Vendor:` field of the output. For example, a DG Clariion disk might return the following Vendor ID:

```
string "DGC".
```

This string must be included in the `vxdump.conf` file as follows:

```
dmp_jbod="DGC" ;
```

If disks from multiple vendors are present in one JBOD disk array, vendor IDs of all vendors should be specified in the `vxdump.conf` file. If this is not done, disks belonging to vendors that are not mentioned in the `dmp_jbod[]` array, are not multipathed by DMP and come under the category of OTHER DISKS. (Other Disks is a separate category of disks not multipathed by DMP and does not come under any of the supported categories in DMP.)

The following example shows how to add entries for multiple vendors in the `dmp_jbod[]` array.

```
dmp_jbod="DGC" ,
        "IBM" ,
        "FUJITSU" ;
```

The specified Vendor ID in the `vxdump.conf` file cannot be more than eight characters long. If a Vendor ID string more than eight characters long is specified in the `vxdump.conf` file, the following message displays on the console:

```
NOTICE: vxvm:vxdump: Invalid Vendor ID = <vendor_id> specified
in vxdump.conf file
```

```
Action: Remove or correct this entry in the vxdump.conf file.
```

4. After the `vxdump.conf` file has been edited to add the required entries, the system should be rebooted so that DMP can multipath the specified JBODs.

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 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.

Using the AP Driver With the 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 Solaris patch 110722-01 (or higher). If you use AP, upgrade AP to 2.3.1 and install Solaris patch 110722-01 before installing/upgrading to this release of VERITAS Volume Manager.

Note The `vxdmp` 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 Solaris patch 110722-01, the installation/upgrade to VxVM 3.2 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 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 3.2 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.

```
# vxdmpadm listctlr all
CTLR-NAME          ENCLR-TYPE          STATE          ENCLR-NAME
=====
c0                 OTHER               ENABLED        others0
pseudo            AP                  ENABLED        ap0
```

```

# vxddmpadm getsubpaths ctrl=mc2
NAME STATE          PATH-TYPE  DMPNODENAME  ENCLR-TYPE  ENCLR-NAME
=====
mc2t35d0s2  ENABLED    -          mc2t35d0s2  AP          ap0
mc2t33d0s2  ENABLED    -          mc2t33d0s2  AP          ap0
mc2t34d0s2  ENABLED    -          mc2t34d0s2  AP          ap0

# vxddmpadm listenclosure all
ENCLR_NAME          ENCLR_TYPE    ENCLR_SNO          STATUS
=====
others0             OTHER OTHER_DISKS          CONNECTED
sena0               SENA          5080020000040750  DISCONNECTED
ap0                 AP            5080020000040750  CONNECTED

# vxddisk list
DEVICE              TYPE          DISK              GROUP            STATUS
=====
c0t0d0s2           sliced       -                 -                offline
c0t1d0s2           sliced       disk02            rootdg           online
mc2t33d0s2         sliced       -                 -                online
mc2t34d0s2         sliced       disk03            rootdg           online
mc2t35d0s2         sliced       disk01            rootdg           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 *ap0*.

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 `mc#t#d#` name
3. Follow the steps in the AP documentation to remove a pathgroup.
4. Run the VxVM configuration daemon, as follows:

```
# vxddctl 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 3.2, ATF/DMP co-existence works *only* with the `CLatf 1.3.1` and `CLatf 3.1.0` package.

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:
 - a. Install ATF.
 - b. Install VxVM.
- ◆ If VxVM is installed with DMP currently enabled, do the following:
 - a. Prevent or suppress DMP.
 - b. Install the ATF driver.
 - c. Install VxVM 3.2, if appropriate.
 - d. If DMP is not re-enabled, re-enable it now.

- ◆ If VxVM is installed with DMP currently prevented or suppressed, do the following:
 - a. Install the ATF driver.
 - b. Install VxVM 3.2, if appropriate.
 - c. If DMP is not allowed or unsuppressed, do this now.

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 A5000, A5100, and A5200 disk arrays. These disk arrays are claimed by VxVM (DMP) under the SEAGATE category.

The A5x00 disk arrays will be claimed by VxVM (DMP) under the SENA category only if the required libraries are present on the system at the time of the installation or upgrade of VxVM.

Please see “[Important Software Patch Information](#)” in the *VERITAS VxVM 3.2 Installation Guide* for information about current required patches.

If the appropriate patches are not on the system at the time of installation or upgrade to this version of VxVM, then the A5x00 disk arrays are claimed in the SEAGATE category. In this case, perform the following steps to claim these arrays in the SENA category.

1. Stop `vxconfigd`
2. Add the appropriate Solaris patch
3. Copy `/usr/lib/liba5k.so.2` to `/etc/vx/slib`
4. Copy `/usr/lib/libg_fc.so.2` to `/etc/vx/slib`



5. Restart vxconfigd

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:        sliced
hostid:      gluon.veritas.com
disk:        name=rootdisk id=933470913.1042.gluon.veritas.com
group:       name=rootdg id=933470911.1025.gluon.veritas.com
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
private:     slice=7 offset=1 len=4711
update:      time=933471058 seqno=0.9
headers:     0 248
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:        sliced
```

```

hostid:      gluon.veritas.com
disk:        name=rootdisk id=933470913.1042.gluon.veritas.com
group:       name=rootdgid id=933470911.1025.gluon.veritas.com
flags:       online ready autoconfig autoimport imported
pubpaths:    block=/dev/vx/dmp/clt33d0s6 char=/dev/vx/rdmp/clt33d0s6
privpaths:   block=/dev/vx/dmp/clt33d0s7 char=/dev/vx/rdmp/clt33d0s7
version:     2.1
iosize:      min=512 (bytes) max=2048 (blocks)
public:      slice=6 offset=1 len=35363560
private:     slice=7 offset=1 len=4711
update:      time=933471058 seqno=0.9
headers:     0 248
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
clt33d0s2    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 ctrlr=/c1
```

NAME	STATE	TYPE	DMPNODENAME	DA-TYPE	DA-SNO
NONAME	DISABLED	-	clt36d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt49d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt38d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt37d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt53d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt34d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt51d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt54d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt52d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt33d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt50d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt32d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt48d0s2	SEAGATE	SEAGATE_DISKS
NONAME	DISABLED	-	clt35d0s2	SEAGATE	SEAGATE_DISKS

Note that when the path used to name the DMP metadvice (that represents a LUN) is removed, DMP uses the name of the alternate path to name the DMP metadvice.



For example, assume that there are two paths to a LUN called `c1t33d0s2` and `c2t33d0s2`. The DMP metadvice that represents the LUN is named `c1t33d0s2` and is in the `ONLINE` state. If the cable connected to controller `c1` is removed, followed by the `vxctl enable` command, the DMP metadvice is renamed to `c2t33d0s2`. As a result, there are two metadvice corresponding to the same LUN visible to the user. Any one of these could be in the `error` state and the other in the `online` state. Note that this is not a problem. However, this results in some error messages when the `vxdisk list <daname>` command is used on the original `daname` that was used to name the DMP device. The Multipathing Information does not display to the user, but the following error message displays:

```
vxvm:vxdisk: ERROR: <da_name> not a valid dmpnode disk access name.
```

The disk paths that are no longer accessible can be removed using the `rm` command.

Minimum Revision Levels for Solaris Patches

The following table specifies the minimum revision level required for each Solaris patch ID.

DISCLAIMER: Patch information is determined at the time of product release. For the most current information, please contact your vendor.

Solaris 2.6	Solaris 7	Solaris 8
105181-28 kernel update	106541-17 kernel update	108528-09 kernel update
105356-18 sd, ssd drivers	106924-06 isp driver	109524-06 ssd driver
105580-16 glm driver	106925-08 glm driver	109657-05 isp driver
	107458-10 dad, sd, ssd, uata drivers	108974-11 sd, uata drivers

Dynamic Reconfiguration

Dynamic Reconfiguration (DR) is a feature available on some high-end SUN Enterprise 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:

```
# vxdmppadm 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 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 board.
4. After the board is replaced (or attached), re-enable the Volume Manager configuration daemon using the following command:

```
# vxdctl enable
```

5. 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, in the figure, above, for attaching System Board 8, run the following commands:

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

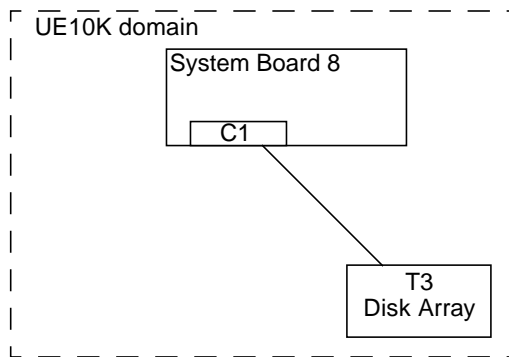
This command insures that the controllers are now available for I/O to the Volume Manager.

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. All disks connected to the board being detached must be removed entirely from VxVM control. Run the command `vxdiskadm` and choose the option:


```
Remove a disk for replacement
```

Run the following command on the disk:

```
# vxdisk rm daname
```
2. Follow the steps in the Sun documentation to detach the system board.
3. 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).
4. Run the Volume Manager configuration daemon using the following command:


```
# vxctl enable
```



- Re-add the disks to Volume Manager using the command `vxdiskadm` and choose the following option:

Replace a failed disk or removed disk

- 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:

- For the Sun StorEdge T3 disk array which is accessible only through System Board 8, follow the procedure for Domain Configuration Type 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 3.2. 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*”.
- ◆ VERITAS Volume Manager volumes configured using devices from the Sun StorEdge A3000/A3500 Array subsystem cannot be part of the root disk group, `rootdg`. Configure StorEdge A3000/A3500 Array devices only to non-`rootdg` disk groups.
- ◆ 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 (see the StorEdge A3000/A3500 Array documentation for the device nodes):

```
/dev/RAID_module01/0s[0-7]
```

and

```
/dev/RAID_module01/0s[0-7].
```

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, if you have an `/etc/vfstab` entry similar to this:

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

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
```
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.

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; however, 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 Array With VxVM

Sun StorEdge T3 Support

This note relates to Sun StorEdge T3 support in two environments:

- ◆ Single host environment
- ◆ Multi host environment

Sun StorEdge T3 Support In A Single Host Environment

The Sun StorEdge T3 should be configured in RW (implicit failover or auto-trespass mode) 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 array.

Sun StorEdge T3 Support In A Multi-host Environment

In a multi-host environment, the Sun StorEdge T3 should be configured in STD (explicit failover or nonauto_trespass) mode. Whenever a host encounters a path failure, VxVM issues an explicit LUN failover command through the alternate path to fail over the ownership on the original path. Then VxVM, on these hosts, will move their active path to their alternate path.

You need to upgrade the Sun StorEdge T3 firmware to 1.17B. The patch for this upgrade is 109115-09.

You should use the `sys mp_support std` command on the Sun StorEdge T3 to enable the STD (explicit failover) mode.

Note The modes cannot be mixed. All hosts in the multi-host environment must use the explicit failover mode.

Working with VxVM Disks on a Sun StorEdge T3

On a T3 or T3+ array with the `c#t#d#` naming scheme, use the following procedures to initialize, remove, uninitialized, or replace VxVM disks.

Note `vxconfigd` is temporarily stopped and restarted each time the mode of the array is changed. Keep in mind that data I/O is always operating during this time.

Initializing a Disk on a T3 or T3+ Array in c##t##d# Mode

1. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming mode to enclosure based format.
2. Issue `vxdisksetup` to add the new disk.
3. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming back to c##t##d# mode.

Removing or Uninitializing a Disk on a T3 or T3+ Array in c##t##d# Mode

1. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming mode to enclosure based format.
2. Issue `vxdg` command to remove the disk from the disk group.
3. Issue `vxdiskunsetup` to remove the disk from VxVM control.
4. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming back to c##t##d# mode.

Replacing a Disk on a T3 or T3+ Array in c##t##d# Mode

1. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming mode to enclosure based format.
2. Issue `vxdg` command to remove the disk from the disk group.
3. Issue `vxdiskunsetup` to remove the disk from VxVM control.
4. Replace the disk.
5. Issue `vxdisksetup` to add the new disk.
6. Use Option 20 of `vxdiskadm` "Change the disk naming scheme" to change the name scheming back to c##t##d# mode.



